

A JOINT COMMUNIQUÉ: THE PSI GANZFELD CONTROVERSY

BY RAY HYMAN AND CHARLES HONORTON

ABSTRACT: Instead of continuing with another round of our debate on the psi ganzfeld experiments, we decided to collaborate on a joint communiqué. The Honorton-Hyman debate emphasized the differences in our positions, many of these being technical in nature. But during a recent discussion, we realized that we possessed similar viewpoints on many issues concerning parapsychological research. This communiqué, then, emphasizes these points of agreement. We agree that there is an overall significant effect in this data base that cannot reasonably be explained by selective reporting or multiple analysis. We continue to differ over the degree to which the effect constitutes evidence for psi, but we agree that the final verdict awaits the outcome of future experiments conducted by a broader range of investigators and according to more stringent standards. We make recommendations about how such experiments should be conducted and reported. Specific recommendations are about randomization, judging and feedback procedures, multiple analysis and statistics, documentation, and the growing role we believe meta-analysis will play in the evaluation of research quality and the assessment of moderating variables. We conclude that psi researchers and their critics share many common goals, and we hope that our joint communiqué will encourage future cooperation to further these goals.

The *Journal of Parapsychology* had planned to publish one more exchange between us on the debate that we initiated in the March 1985 issue (Honorton, 1985; Hyman, 1985). In fact, one of us had already written and submitted his reply, and the other was preparing a rejoinder when we encountered each other at the 1986 meetings of the Parapsychological Association. The idea of replacing another round of exchanges with this joint communiqué emerged from a discussion during a luncheon meeting.^{1,2} During the discussion we realized that each of us had not fully and accurately understood the other's position on some of the major issues dividing us. In addition, much of our disagreement at this stage involves technicalities and differences of opinion about the proper ways to assign

¹Marcello Truzzi participated in this discussion. We would like to thank him for his encouragement and suggestions.

²To the best of our knowledge, this is the first time a parapsychologist and a critic have collaborated on a joint statement of this type. Hyman prepared the first draft and we continued exchanging drafts until we had one we were both satisfied with. For those who are interested, the final product is draft 4.

and rate studies on specific attributes. To put emphasis on these details detracts from the broader and more important propositions, on which we find ourselves in agreement.

These propositions relate in general to how psi researchers and critics can work together toward the resolution of their differences. Specifically, they relate to how we believe psi ganzfeld experiments should be conducted and reported in the future.

GENERAL AREAS OF AGREEMENT

As in any other area of scientific inquiry, research in parapsychology requires continual scrutiny and criticism. Both critics and parapsychologists want parapsychological research to be conducted according to the best possible standards. The critic can contribute to this need only if his criticisms are informed, relevant, and responsible.

As to the psi ganzfeld data base, we agree, as our earlier exchanges indicate (Honorton, 1983, 1985; Hyman, 1983, 1985), that the experiments as a group departed from ideal standards on aspects such as multiple testing, randomization of targets, controlling for sensory leakage, application of statistical tests, and documentation. Although we probably still differ about the extent and seriousness of these departures, we agree that future psi ganzfeld experiments should be conducted in accordance with these ideals. In the second section of this joint communiqué, we shall make a number of specific recommendations about the conduct and documentation of future psi ganzfeld studies. It is our hope that these recommendations will lay the groundwork for a new round of studies that will serve to resolve the differences remaining between us.

Although we probably still differ on the magnitude of the biases contributed by multiple testing, retrospective experiments, and the file-drawer problem, we agree that the overall significance observed in these studies cannot reasonably be explained by these selective factors. Something beyond selective reporting or inflated significance levels seems to be producing the nonchance outcomes. Moreover, we agree that the significant outcomes have been produced by a number of different investigators.

Whereas we continue to differ over the degree to which the current ganzfeld data base contributes evidence for psi,³ we agree that

³As regular readers of the *Journal* know, the term *psi phenomena* was introduced

the final verdict awaits the outcome of future psi ganzfeld experiments—ones conducted by a broader range of investigators and according to more stringent standards.

The strongest disagreements between us might appear to be over the relationship in the data base between “flaws” and study outcome. Honorton finds no significant correlation between indices of study quality and study outcome. Hyman agrees that there is no significant correlation between study outcome and some procedural indicators, such as multiple analysis, sensory leakage, statistics, and security. But he finds a positive correlation between study outcome and other procedural indicators, such as suboptimal randomization, feedback, and inadequate documentation.

Which correlation one obtains depends on how the “flaws” are assigned to individual studies, how one orders the seriousness of flaws when constructing scales, how many different attributes are included as flaws, and similar judgments. But these differences, no matter how controversial, should not be allowed to obscure our agreement that the present data base does not support any firm conclusion about the relationship between flaws and study outcome. Our disagreements about the actual correlation only emphasize this point.

If psi is responsible for the outcomes obtained in this data base, then the ganzfeld experiment should continue to produce successful outcomes when the various problems that Hyman pointed out are eliminated. Indeed, what differentiates the ganzfeld debate from many earlier controversies between psi researchers and critics is that the claim is one of replicability. Consequently, the best way to resolve the controversy between us is to await the outcome of future ganzfeld psi experiments. These experiments, ideally, will be carried out in such a way as to circumvent the file-drawer problem, problems of multiple analysis, and the various defects in randomization,

by Thouless and Wiesner as a neutral label denoting unexplained interactions between organisms and their environment. Their intention was to avoid the surplus meaning associated with terms such as *extrasensory perception*. Although this usage is generally understood within parapsychology, as reflected in the glossary definitions in this *Journal*, the term *psi* has unfortunately taken on broader connotations within the popular culture. Even within parapsychology, the terms *psi* and *paranormal* are sometimes used interchangeably and in a way that confuses description with explanation.

Consistent with the original usage, the term *psi* in this paper simply denotes a communications anomaly. No particular explanation of the anomaly is intended, nor do we believe any is warranted at the present time. We suggest that communication between parapsychologists and other scientists could be improved if this distinction were maintained.

statistical application, and documentation pointed out by Hyman. If a variety of parapsychologists and other investigators continue to obtain significant results under these conditions, then the existence of a genuine communications anomaly will have been demonstrated. The demonstration of an anomaly, of course, does not explain it. Such a demonstration would, however, be very important because it would require acknowledgment that there is, indeed, something to be explained, and the debate would then shift toward such efforts. Whether the anomaly is ultimately to be considered "paranormal" will, as Palmer (1986) suggests, depend on further developments such as the extent to which the findings can be brought under lawful control and the construction of a positive theory of the paranormal.

On the other hand, if the findings can all be attributed to various artifacts, this too is important to determine. Discovering the nature of such artifacts and how they are produced could have important methodological implications for all scientific inquiry. Thus, we agree that further research in this area is important, not only for parapsychology, but for science generally. And we believe it is essential, in order to develop a clear picture of what is actually going on, that the research should be conducted not only by parapsychologists, but by a range of investigators with diverse opinions concerning psi. Studies conducted by investigators who are skeptical of the psi hypothesis would be particularly useful from a number of perspectives. It is possible, for instance, that such studies might reveal potential sources of artifact that have been overlooked or that are not obvious from analysis of existing research reports. Further, investigators favoring conventional explanations of parapsychological findings could contribute substantially to a resolution of the psi controversy by systematically testing and delimiting the explanatory power of various proposed alternative hypotheses (Palmer, 1986).

Finally, before moving on to our recommendations for future psi ganzfeld studies, we believe it is appropriate to say a few words about the process in which we are engaged. As is evident from what has been said above, there are areas in which we continue to disagree. We agree to disagree. Even though our continuing disagreements about the degree to which the existing studies in this area contribute evidence for psi, for example, are not inconsequential, we fully respect each other's position and we disassociate ourselves from the more strident advocates on both sides of the psi controversy who would label those with opposing views by such pejoratives as "prejudiced," "credulous," and "irrational."

RECOMMENDATIONS FOR FUTURE PSI GANZFELD EXPERIMENTS

Although much of what we say in this section might also apply to other areas of parapsychological research, we will make our recommendations specific to the ganzfeld psi experiment and its data base that we discussed previously (Honorton, 1985; Hyman, 1985). The recommendations are intended to illustrate what a parapsychologist and a critic might accomplish when trying to seek common grounds for agreement. What follows are specific recommendations to parapsychologists and other investigators who intend to conduct a ganzfeld psi experiment.

Control for Sensory Leakage

We agree that future investigators should strive to eliminate all possibilities for sensory communication between sender and receiver—both during the ganzfeld session and at judging. The typical two-experimenter psi ganzfeld experiment effectively eliminates sensory leakage during the actual ganzfeld period. The use of duplicate target pools or the binary coding system (Honorton, 1975) guards against sensory leakage at the time of judging. Proper attention to monitoring and recording the actual target should undermine the possibility, suggested by Hyman, of leakage from receiver to sender during feedback. Fortunately, the use of single target pools has disappeared in recent ganzfeld psi experiments.

Randomization of Targets

We agree that more careful attention needs to be given to the procedures for selecting targets and that the procedures should be thoroughly documented.

The method of target selection should be described in full. The following details should be included: (a) the person performing the randomization, (b) the specific source of randomness, (c) the method of sampling the random source (i.e., obtaining entry points for random number tables, seeds for pseudorandom generators, or specific values for hardware random generators). To illustrate what we agree would constitute adequate documentation of randomization procedures for each of the above random sources, consider a hypothetical ganzfeld experiment involving 20 target pools of four pictures each:

Target preparation was performed by R.H., a member of the laboratory staff who was not otherwise involved in the experiment. . .

Random number tables. Pools and targets were selected using the RAND tables (RAND Corporation, 1955). An entry point into the RAND table was obtained for each session as follows. The first digit of the row was determined by coin toss ("heads" = 0, "tails" = 1). A deck of numbered cards (0-9) was then shuffled and cut four times. The uppermost card for each iteration provided subsequent digits for the row. The block (0-9) was then determined by again shuffling the numbered deck. Cards bearing the digits 6-9 and 0 were then removed from the deck, which was again shuffled to determine the specific column (1-5) within the block. The first two digits within the range 01-20 thus provided the pool for the session. R. H. then removed the appropriate judging pool and left it where it could be retrieved by the experimenter. Only after this was done was the actual target for the session determined. This was done by repeating the above procedure and obtaining the first digit within the range 1-4.

Pseudorandom generators. Pools and targets were selected using the random number generator function in the Applied Statistics module of a Texas Instruments TI 59 Programmable Calculator (Texas Instruments, 1977). The seed was obtained by subtracting the six digits comprising the subject's birthdate from the six digits comprising the date of the session. A uniform random number within the range of 1-20 provided the pool for the session. R.H. then removed the appropriate judging pool and left it where it could be retrieved by the experimenter. Only after this was done was the actual target for the session determined. The next random digit within the range of 1-4 was the target.

Hardware random number generators. The targets were selected using a PsiLab II random number generator interfaced to an Apple II computer (Psychophysical Research Laboratories, 1984). A BASIC program sampled the RNG, such that the first byte value returned within the range of 1-20 was the pool. R.H. then removed the appropriate judging pool and left it where it could be retrieved by the experimenter. Only after this was done was the actual target for the session determined. This was done by repeating the above procedure and obtaining the first digit within the range of 1-4.

Although random number generators are often today more convenient than tables of random numbers, several caveats are in order regarding their use in serious research applications. In general, we do not recommend use of microcomputer random functions. The algorithms used are generally not documented, and some have been shown to produce spurious results (e.g., Hansen, 1986). A good discussion of the characteristics of some of the more widely used

pseudorandom algorithms is given in Radin (1985). Hardware random number generators can have design flaws or may develop intermittent problems that will lead them to fail. And even though a brief description of the circuitry and its theory of operation is desirable, technical descriptions of the hardware cannot be substituted for empirical tests of the output. Ideally, randomness tests would be done on the actual targets used in an experimental series. However, owing to the typically small sample size of psi ganzfeld studies, such tests would be of little value. Control tests should be reported to insure adequate randomness of the targets. Because ganzfeld experiments involve only one target selection per session, sequential bias is not likely to be an issue as it could be in other areas of psi research, and the ganzfeld investigator can restrict his or her attention to a frequency analysis allowing assessment of the degree to which targets occur with equal probability. A good discussion of randomness tests is given by Davis and Akers (1974).

Hyman believes that the best way to assure adequate randomization is to include empirical, in addition to the usual theoretical, baselines. One way to do this would be to systematically compare the percipient's first choice both against the intended target and against the intended target for a control trial (which could be the actual trial for another percipient). If the randomization procedure is adequate, the control comparisons should produce observed means and standard deviations consistent with the theoretical distribution on the null hypothesis. This recommendation is identical to the "cross-check" method used in the early card-guessing experiments (Rhine et al., 1940/1966, p. 46).

Judging and Feedback

We agree that the judging and feedback procedures should be presented in greater detail than has generally been the practice in past ganzfeld studies. Specifically, the report should explicitly document the following procedures: (a) the manner in which persons knowing the identity of the actual target (i.e., the sender and sender's experimenter) remain isolated from the receiver and receiver's experimenter until completion of judging; (b) the instructions given to the receiver for judging; (c) how the judging pool is presented to the subject; (d) the manner in which the subject's ranks or ratings are recorded; and (e) how feedback to the actual target is delivered at completion of the subject's judging.

Multiple Analysis

The problem of multiple analysis pervades all the sciences. Determining the size of the total critical region is often difficult even when the investigator has conscientiously set out in advance the tests that will be made. More typically the investigator has a more or less general idea of the hypotheses to be tested, but the precise details have not been worked out in advance. Under these conditions, the precise indices, cutting points, and tests are constructed after the data have been collected and assembled. The temptation is strong to tailor the specifics of the testing to the peculiarities of the data.

Even though it is not possible to make a generalized recommendation that will meet all contingencies, clearly investigators should specify all the confirmatory tests, as well as the precise critical region in advance of collecting the data, and such specification of confirmatory tests should be explicitly stated in the experimental report. Adherence to this recommendation should not be taken as being inconsistent with exploratory data analysis. The point of the recommendation is to prevent confusion between confirmatory tests and suggestive findings that require confirmation by future experiments. When multiple tests are planned, appropriate adjustments should be made to keep the total overall error rate within the commonly accepted region. One approach involves using the Bonferroni inequality (e.g., Rosenthal & Rubin, 1984).

We recommend that future investigators consider the possibility of increasing statistical power by using, with appropriate adjustments, two or more of the several indices that have been used as indicators in psi ganzfeld experiments.

Hyman believes either theoretical or empirical investigations might suggest that a linear combination of two or more of these indices could usefully increase statistical power. Or, it might turn out that more power can be achieved by performing separate tests on two or more of these indicators and then adjusting the overall level of significance appropriately. It is not clear at this time what might be the best combination. If the investigator decides to use just two indices, for example, statistical considerations might suggest choosing those two that are least correlated. This would argue for using both direct hits and binary hits. Simulation studies by Hyman indicate that these two measures correlate 0.61, whereas the inter-correlations between any other pair of the most common indicators are approximately 0.80 or higher. On the other hand, the studies in the current data base that used binary hits rather than direct hits

appeared to yield less impressive results. Those studies that used the special slide pool allowing use of Honorton's binary coding system (Honorton, 1975) indicate that the resulting index is uncorrelated with the four major indices of direct hits, binary hits, sum of ranks, and normalized ratings. However, the binary coding index also seems to yield smaller and less significant effects (Hyman, 1985).

Honorton believes a good case can be made for using both direct hits and sum of ranks measures. Because the two measures are discretely distributed, the penalty required for using both is minimal. Consider, for example, a study involving 20 trials and a hit probability of .25. If alpha is set to .05 and the direct hits measure alone is used, a significant outcome will be achieved with 9 or more direct hits ($p = .041$). If the sum of ranks measure alone is used, significance will be achieved with a sum of ranks equal to or less than 41 ($p = .036$). If both indices, adjusted by the Bonferroni method, are used, significance will be achieved by either 10 direct hits or a sum of ranks equal to 40. Thus, the added flexibility achieved by allowing use of either measure is, in this case, purchased at a cost amounting to one additional direct hit. The Bonferroni method is overly conservative. As Hyman's simulation shows, direct hits and sum of ranks are highly correlated. Hansen (1986) has recently reported a simulation study involving these two measures and is preparing tables for various sample sizes that will provide more accurate p levels for the use of both direct hits and sum of ranks.

File-Drawer and Retrospective Experiments

Given the Parapsychological Association's policy of actively discouraging the selective reporting of "positive" results, the file-drawer problem is probably less acute in parapsychology than it is in many other scientific disciplines. Certainly reports of nonsignificant outcomes are far more common in the parapsychological literature than in other areas of psychology. This is not to imply, however, that the file-drawer problem is nonexistent in parapsychology. Investigators should bear in mind that registration of statistically nonsignificant outcomes is essential to the development of a realistic appraisal of a research area and that a study's value is independent of its statistical significance.

As to Hyman's suggestion concerning "retrospective" experiments, we recommend that along with specifying the critical region, the investigator should also specify in advance the status of the experiment. Designations such as *classroom exercise*, *confirmatory experi-*

ment, or process-oriented will help future reviewers to both classify and properly evaluate the results.

Statistics

Over 20% of the experiments in the meta-analytic sample of 28 studies contained errors in the use of statistical tests. Although some of these errors may not have had serious consequences, their existence should be a cause for concern to the parapsychological community. We believe that the parapsychological journals, along with the authors, share responsibility for insuring the adequacy of statistical tests used in empirical contributions and that some of these problems could be avoided if authors adhered to the following recommendations printed on the inside back cover of this *Journal*:

1. State concisely the precise statistical formulation of the hypothesis being tested and list it in advance of the results section. It is recommended that the type of statistical tests that are planned be given along with the hypothesis.
2. For any statistical analysis that was not preplanned, give a brief statement of why it was done; the probability value should be placed close enough to this statement that its association is obvious.
3. When statistical analyses are done, report not only the inferential statistics (e.g., *t* values) but also the descriptive statistics for the data evaluated (e.g., group means and standard deviations). Also, report the actual values of correlation coefficients, not simply that a correlation is significant or nonsignificant.
4. Have the data and statistical analyses independently rechecked before submitting the paper.

Documentation

In general, we believe that readers (including research analysts and prospective replicators) should be able to reconstruct the author's procedures from the descriptions provided in the experimental report. Although this is not common practice in science generally, we believe it is important in areas such as parapsychology where routine replicability cannot be taken for granted. More detailed exposition of methods and procedures should serve not only to aid evaluation of research quality, but also to increase the likelihood that other investigators will be able to replicate the original investigator's results successfully. As for future ganzfeld psi experiments, we recommend that, in addition to the procedural details de-

scribed above, investigators routinely supply information on the following: the training, supervision, and qualifications of student experimenters; information on the subject population, including sources of subject recruitment and prior psi-testing experience; the individual ranks and target selections; the acquaintanceship of sender and receiver; the status of the experiment (confirmatory, exploratory, exercise, etc.); and similar information that is germane to the evaluation and replication of the study.

The Role of Meta-Analysis

The standards and recommendations we have discussed so far, for the most part, apply to the individual experiment. Indeed, almost all the guidelines for doing good research are aimed at the individual experiment. And the statistical procedures have been developed and taught with the idea that they apply to the evaluation of a single experiment. But scientific inquiry is cumulative and the outcome of a single experiment rarely, if ever, determines the acceptance or rejection of laws and theories. Science progresses by the cumulative outcomes of many experiments done by many investigators.

This fact has been recently recognized in the contemporary interest in the development of formal techniques for the statistical integration of a series of experiments. The field of meta-analysis is still in its infancy and somewhat controversial. Some of the controversy, as reflected in our debate, deals with the extent to which meta-analysis can compensate for the individual inadequacies of the specific experiments that are included in the data base.

Nevertheless, meta-analysis realistically emphasizes that scientific evidence rests on the consistency of results across many experiments. Before the focus on meta-analysis, the individual investigator designed, conducted, and reported the results with little, if any, consideration of how this particular experiment fit into a larger series of experiments.

Our next recommendation takes into account the growing role we believe meta-analysis will play, both in the evaluation of research quality and in the assessment of moderating variables. We urge parapsychological investigators to plan and report their experiments with the idea that their single experiment will contribute to a future meta-analysis. Much of this information could be encapsulated in summary tables at the end of the research report, as illustrated by Table 1.

TABLE 1
ILLUSTRATIVE STUDY SUMMARY

Session	Subject information									
	ID	Recru	Belief	Pract	Test	Gz	Acq	Tgt	Resp	Rank
1	1	AD	7	1	1	0	0	3B	3D	2
2	2	ST	2	0	0	0	0	11A	11A	1
3	3	OS	5	1	0	0	0	6C	6B	3
4	4	EA	3	0	1	1	1	1D	1D	1
5	5	VI	6	1	1	0	2	17C	17A	4
.
.
.
.

Note. Abbreviations are defined as follows:

ID = subject ID.

Recru = source of subject recruitment.

AD = response to newspaper ad.

ST = student volunteer.

OS = recruited by other subjects.

EA = acquaintance or friend of experimenter.

VI = laboratory visitor.

Belief = belief in psi (1 = low, 7 = high).

Pract = practices such as meditation (0 = no, 1 = yes).

Test = prior psi-testing experience; not ganzfeld (0 = no, 1 = yes).

Gz = prior psi ganzfeld experience.

Acq = sender/receiver acquaintance.

0 = none; sender is laboratory staff member not previously acquainted with subject.

1 = lab friend; subject is a friend/acquaintance of laboratory staff member serving as sender.

2 = friend; sender is a friend/acquaintance of the subject.

Tgt = pool and target for session (e.g., 3A = Pool 3, element 'A').

Resp = pool element selected as first choice by subject.

CONCLUSION

In making these recommendations, we recognize the need to distinguish between ideals and practicality. We believe that the above recommendations are consistent with what is realistically attainable given the current resources of parapsychology. The psi ganzfeld paradigm is now over a decade old, and though the need for innovation and individual creativity is as great as ever, we believe there is also a need for greater discipline and standardization. We

hope that parapsychological investigators and journal editors will welcome our suggestions and view them as a constructive step forward, one that, with their active cooperation, could lead to a broader-based consensus on at least the basic empirical "facts."

Our final recommendation, unlike the others, is probably not feasible under present circumstances, but it is proposed here only to indicate how many of the problems under debate could be avoided if certain ideals could be achieved.

Many of the problems we encountered in evaluating the ganzfeld psi experiments could be avoided in future experiments if the reviewers could be sure that they were dealing with the entire population of relevant studies and could insure the internal validity of those studies. Ideally, the best way to achieve this would be to sponsor a systematic replication series under the auspices of a neutral agency such as the National Science Foundation. The sponsoring agency would establish guidelines and rules based on the joint recommendations of successful investigators and knowledgeable critics. The guidelines would then delimit the experimental designs, the investigator-base, and the time frame for the experimental series, as well as the basic framework for a subsequent meta-analysis of the series as a whole.

The writing of this joint communiqué convinces us that, despite obvious differences, parapsychologists and their critics share many common objectives. These commonalities rarely are noticed in the debates, which focus on the differences. Yet such commonalities hold the key for how the parapsychologist and the critic can join forces to achieve the ends to which they both aspire.

REFERENCES

- DAVIS, J. W., & AKERS, C. (1974). Randomization and tests for randomness. *Journal of Parapsychology*, **38**, 393-407.
- HANSEN, G. (1986). *Monte Carlo methods in psi research*. Paper presented at the 29th Annual Convention of the Parapsychological Association.
- HONORTON, C. (1975). Objective determination of information rate in psi tasks with pictorial stimuli. *Journal of the American Society for Psychical Research*, **69**, 353-359.
- HONORTON, C. (1983). Response to Hyman's critique of psi ganzfeld studies. In W. G. Roll, J. Beloff, & R. A. White (Eds.), *Research in parapsychology 1982* (pp. 23-26). Metuchen, NJ: Scarecrow Press.
- HONORTON, C. (1985). Meta-analysis of psi ganzfeld research: A response to Hyman. *Journal of Parapsychology*, **49**, 51-91.
- HYMAN, R. (1983). Does the ganzfeld experiment answer the critics' objections? In W. G. Roll, J. Beloff, & R. A. White (Eds.), *Research in parapsychology 1982* (pp. 21-23). Metuchen, NJ: Scarecrow Press.

- HYMAN, R. (1985). The ganzfeld psi experiment: A critical appraisal. *Journal of Parapsychology*, **49**, 3-49.
- PALMER, J. (1986). Progressive skepticism: A critical approach to the psi controversy. *Journal of Parapsychology*, **50**, 29-42.
- PSYCHOPHYSICAL RESEARCH LABORATORIES. (1984). *PsiLab // user's manual*. Princeton, NJ: Psychophysical Research Laboratories.
- RADIN, D. (1985). Pseudorandom number generators in psi research. *Journal of Parapsychology*, **49**, 303-328.
- RAND CORPORATION. (1955). *A million random digits with 100,000 normal deviates*. New York: The Free Press.
- RHINE, J. B., PRATT, J. G., SMITH, B. M., STUART, C. E., & GREENWOOD, J. A. (1940/1966). *Extra-sensory perception after sixty years*. Boston: Bruce Humphries.
- ROSENTHAL, R., & RUBIN, D. B. (1984). Multiple contrasts and ordered Bonferroni procedures. *Journal of Educational Psychology*, **76**, 1028-1034.
- TEXAS INSTRUMENTS. (1977). *Applied statistics*. Lubbock, Texas: Texas Instruments Inc.

Department of Psychology
University of Oregon
Eugene, OR 97403

Psychophysical Research Laboratories
301 College Road East
Princeton, NJ 08540