

Experiment One of the SAIC Remote Viewing Program: A critical re-evaluation

Dr Richard Wiseman
University of Hertfordshire
and
Dr Julie Milton
University of Edinburgh

Abstract

The Central Intelligence Agency recently commissioned a report on the effectiveness of US government funded research in demonstrating the existence of a remote viewing effect that could be used for intelligence-gathering purposes. The evaluators focused their attention on successful studies recently conducted by Science Applications International Corporation (SAIC) and concluded that the studies contained no discernible flaws. This paper first outlines xx a key study in the report (referred to as Experiment One) and then examines whether the study contained four potential methodological problems. The paper demonstrates that there are severe problems associated with reconstructing important aspects of the study and that two of the flaws may be able to account for the study's outcome.

Introduction

For over 20 years the US government has funded experiments examining the possible existence of 'remote viewing' -- the ability to psychically acquire information from a distant location. During the early 1970s this work was carried out at the Stanford Research Institute (SRI). The program was transferred to SRI International (formerly SRI) in 1973 and then continued between 1992 and 1994 at Science Applications International Corporation (SAIC).

In September 1995, the American Institutes for Research (AIR), contracted by the Central Intelligence Agency, assembled a 'blue ribbon' panel to evaluate this research (Mumford, Rose & Goslin, 1995). The panel included two reviewers chosen for their expertise in parapsychological research, namely Dr Jessica Utts, a professor of statistics at the University of California at Davis and Dr Ray Hyman, a psychologist at the University of Oregon. Dr Michael Mumford and Dr Andrew Rose, both of them senior behavioural scientists and experts in research methods at AIR, together with Dr Lincoln Moses, a professor of statistics at Stanford University and Dr David Goslin, President of AIR and review coordinator, completed the panel. The resulting AIR report contained two major reviews of the research (Utts, 1995a; Hyman, 1995a) and a concluding section that outlined the main points of agreement and disagreement between the two reviews. Both Utts' and Hyman's reports have since been reproduced in two journals (Utts, 1995b, 1996; Hyman, 1995b, 1996).

Utts (1995a) and Hyman (1995a) directed their attention mostly to the SAIC experiments, in part because, unlike in the case of SRI work, they had access to written details of the full set of ten SAIC studies. Dr Edwin May, SAIC's project supervisor for the remote viewing work, was available to provide details that were not covered in the written reports. Utts took the approach that if the SAIC experiments' successful outcomes had been obtained under methodologically rigorous conditions, then it would be reasonable to conclude that there was evidence for the existence of 'anomalous cognition' or ESP. One of the most important parts of her report therefore involved assessing the safeguards employed in the SAIC remote viewing studies by means of a checklist of eight specific issues. These issues, partly derived from criticism of some of the early SRI experiments [Footnote 1], included concerns about using an adequate randomisation source, preplanning the statistical test for the study's outcome, and so on. For the only detailed assessment of an individual study in her report, Utts selected a single experiment from the group as an example of her evaluation method for all of the studies and described how its methodology contained adequate safeguards against the eight potential methodological loopholes. The study chosen was Experiment One, which had obtained statistically significant results. Both Hyman and Utts agreed that this experiment didn't contain any obvious flaws.

The publication of the AIR report generated a huge amount of media interest, much of it favourable to the view that the SAIC experiments supported the existence of ESP. Coverage in the US included items on ABC News Nightline, CNN Prime News, CNN Larry King Live, CBC World News and articles in USA Today, Washington Post, International Herald Tribune, and news of report appeared in the media of other countries including the UK, India, Germany and China.

The first author (henceforth referred to as RW) has an interest in assessing strong claims for the existence of psi and the AIR report constituted a strong claim, indirectly by the US government, that the SAIC studies were methodologically 'clean' and yet still successful. RW therefore became interested in attempting a replication and so in February 1996 contacted May regarding this plan. May was enthusiastic about the idea and asked RW to propose an experimental protocol. Before doing so RW thought it would be useful to examine the protocol used in Experiment One. He chose this study for three reasons. First, it had been discussed in detail and endorsed by Utts (1995a) in her AIR report. Second, it was one of the few SAIC studies that had been published in a peer-reviewed journal (Lantz, Luke & May, 1994) rather than as an internal report. Third, like all of the SAIC experiments, its design had been examined before the study began by a multidisciplinary 'Scientific Oversight Committee' that included a Nobel Prize winning physicist and internationally known professors of statistics and psychology. As such, RW expected the experiment's methodology to be a suitable template for the proposed replication.

Because the information in Lantz et al.'s (1994) paper and the AIR report was not sufficient to fully evaluate the experiment's protocol, RW asked May for additional details. May kindly sent a copy of an SAIC technical report about the experiment (May & Lantz, 1991) and provided additional information via correspondence and meetings with RW.

However, as the assessment of the protocol progressed it became apparent that Experiment One's methodology contained some problems. RW therefore decided to examine the

protocol in detail, both to reassess the validity of the experiment and to help ensure that future studies, including his own replication attempt, would avoid similar pitfalls. May agreed to provide the necessary unpublished information concerning expt Ones procedure, in order that RW could conduct a detailed critique for publication. The outcome of this critical re-evaluation is presented here.

Summary of Experiment One

Lantz et al. (1994) describe three experimental roles for SAIC personnel in Experiment One, without identifying the holders of those roles by name. May (personal communication to RW, 29 April, 1996) informed RW that the Principal Investigator was Nevin Lantz, the Experiment Co-ordinator was Wanda Luke, and the Analyst (independent judge) was May himself. To clarify later discussion, these experimenters are referred to by name instead of by title in the following account of Experiment One summarised from Lantz et al.

Experiment One was carried out in 1992 and employed a 2x2 design to explore whether receivers' psychic performance was affected by (i) the presence or absence of a sender and/or (ii) the type of target used during a trial. Half of the trials therefore involved a sender while the other half had no sender, and half of the trials involved 'static' targets (photographs from National Geographic magazine) while the other half used 'dynamic' targets (a varied selection of video clips). Prior to the experiment the targets had been grouped into sets of five, each set composed of only static targets or only dynamic targets. Targets within any one set were chosen to be as different from one another as possible.

Five receivers were involved in the experiment, each doing 10 trials in each of the 4 conditions. Two of the viewers lived in California and the others resided in Kansas, New York and Virginia, respectively. Lantz, who acted as sender, was in Lititz, Pennsylvania. May and Luke were at SAIC in Menlo Park, California.

Prior to the study, Luke randomly selected 40 targets (20 static, 20 dynamic) for each participant. A copy of each target was placed into an envelope and a trial number (1 to 40)

written on the outside of the envelope. Half of the static targets and half of the dynamic targets were randomly assigned to the 'no-sender' condition and the envelopes containing these targets were sealed. The remaining targets were assigned to the 'sender' condition and these envelopes were left unsealed.

Luke then mailed all of these targets to Lantz. For each trial Lantz selected the appropriate envelope at a prearranged time. For 'no-sender' trials, he simply placed the unopened envelope on his desk. For 'sender' trials, he looked at the target for approximately 15 minutes. During this time the receiver, at his or her home, produced a 'response' by writing and/or drawing his or her impressions of the target. At the end of the trial the receiver faxed this response to Lantz. Lantz then sent the target to the receiver by return mail, to act as feedback. The receiver subsequently sent the original copy of the response and the copy of the target to Luke at SAIC.

When Luke received the response, she removed the receiver's name and the date and time of the trial. She then gave May, the independent judge, the receiver's response and the appropriate 'target set' (i.e., the set containing a copy of the actual target and four decoys). He ranked the target set items in order of their correspondence to the receiver's response, giving a rank of 1 to the target most similar to the response.

Using Solfvin, Kelly and Burdick's (1978) sum of ranks statistic [Footnote 2] to assess the outcome, we find that the whole study was statistically significantly above chance ($N=200$, effect size, $Z/N^{1/2}=12$, $p<.043$, 1-tailed). Lantz et al. (1994) also found that trials involving the static targets were independently significant ($N=100$, effect size $=.24$, $p<.0073$) whereas those with dynamic targets were at chance ($N=100$, effect size $=.00$, $p=.5$). There was no significant difference in scoring between those trials that involved a sender and those that did not, nor were there any significant interactions between the sender and target variables.

A critical re-evaluation

Examination of Lantz et al.'s (1994) description of the study indicated the possibility of four methodological loopholes. The following sections describe our attempts to clarify to what extent each might have affected the study's outcome. We are indebted to May for generously taking the time to address the issues raised.

(i) Safeguards against the experimenters inadvertently altering the receivers' responses

In ESP research it is generally accepted that any individual who records, transcribes or edits a receiver's responses should be blind to the identity of the target (Milton & Wiseman, 1997). This safeguard prevents experimenters inadvertently biasing the response to match the target, for example by unconsciously tending to edit out inaccurate sections of the response more than accurate sections. In Experiment One, RW noticed what appeared to be an opportunity for a similar potential flaw, namely, for a non-blind experimenter (Luke) to accidentally leave handling cues on the receivers' responses that might unconsciously bias the independent judge.

According to Lantz et al.'s (1994) account, Luke had two opportunities to discover the identity of each trial's target. First, she randomly selected the targets, packaged them and shipped them to Lantz. Second, the receivers returned the feedback copy of the target to Luke. Lantz et al. also note that the receivers sent their responses to Luke, who removed the receiver's name, the date and trial number from the response sheet. In short, Lantz et al.'s account of Experiment One described no safeguards against Luke knowing the target for each trial and accidentally leaving marks (e.g., smudges) on the response sheets that might have cued the independent judge to the target's identity without him having consciously observed the cues.

Given that this procedure outlined by Lantz et al. (1994) seemed to contain a potential flaw, RW asked May whether their description of this aspect of the procedure was correct. May (personal correspondence, 29 April 1996 [Footnote 3]) confirmed that there had only been one person responsible for selecting the targets and removing the trial-identifying information from the responses when they were returned to SAIC, and that that person was

Luke. He also stated that the receivers 'never sent the target material back' to SAIC (personal correspondence, 29 April 1996). However, Lantz et al. state that 'The target copy and original response were subsequently sent to the Experiment Coordinator [Luke] in Menlo Park, CA [SAIC]' (p. 289). May's memory of this aspect of the protocol conflicted with the account in Lantz et al. (see the first two rows of Table 1). RW asked May about this discrepancy and he replied that he could not have remembered it correctly and that either the Lantz paper or his memory was incorrect (personal correspondence, 29 April 1996).

May then asked Luke and Lantz about this aspect of the protocol and sent the resulting 'collective memory' to RW (personal correspondence, 29 August 1996). This collective memory differed in two important ways from the information in Lantz et al. (1994) and the situation originally recalled by May (see the first three rows of Table 1). Both Lantz et al. and May stated that only Luke selected the targets, whereas the collective memory indicated that three people were involved. Also, Lantz et al. and May stated that Luke received and edited the responses whereas the collective memory stated that DeGraff carried out these tasks.

May then checked again with the other SAIC personnel and sent RW a second collective memory which differs from the information presented by Lantz et al. (1994), May and the first collective memory (see the third row of Table 1). First, the second collective memory stated that Macgowan and DeGraff initially selected and packaged the targets. Lantz et al. and May stated that these tasks were carried out by Luke, whereas the first collective memory stated that all three individuals were involved. Second, the second collective memory stated that the targets were returned to Macgowan and DeGraff. In contrast, Lantz et al. stated that they were returned to Luke whereas May didn't think they were returned at all. Finally, according to the second collective memory the responses were returned to Luke. This is in agreement with Lantz et al. and May, but not the first collective memory, which stated that they were returned to DeGraff.

In response to the PA paper xx May then sent RW a final version of events (personal correspondence, 30 April 1997). He noted that this version was the most reliable as he had spent a considerable amount of time carefully interviewing the people involved and had gone to the trouble of recording these interviews on videotape. In this final version of events the targets were selected by DeGraff, McGowan and a student assistant. Luke received the responses and was told the target pack from which the target was chosen, but was blind to the target. The final row in Table 1 shows the differences between this version of events and the previous versions.

	Target selection	Handling returned targets	Handling and editing responses
Lantz et al.	Luke	Luke	Luke
May	Luke	Targets not returned	Luke
First collective memory	Luke, DeGraff, Macgowan	Not mentioned	DeGraff
Second collective memory	DeGraff, Macgowan	DeGraff, Macgowan	Luke
Third collective memory	DeGraff, Macgowan Student assistant	Not mentioned	Luke

Table 1: Differences in the accounts of Lantz et al. (1994), May's first recollection, the experimenters' first, second and third collective memory regarding who selected the targets, and handled the targets and responses at SAIC in Experiment One.

To summarise, from Lantz et al.'s (1994) account of Experiment One it appeared possible that a non-blind experimenter may have handled receivers' responses before passing them on

to an independent judge, allowing an opportunity for inadvertent sensory cueing. xx May produced several different accounts of this aspect of the experimental protocol and so it is difficult to know the degree of certainty that one should attach to any of the accounts. However, assuming that the final version is correct, the possibility that cues could have been inadvertently given to the independent judge in the way suggested is ruled out.

(ii) Safeguards against receivers altering their responses after finding out the identity of the target

Several commentators on parapsychological methodology have recommended that the target should remain shielded from the receiver until after his or her response has been secured (Akers, 1984; Morris, 1986). This safeguard is important not only to rule out deliberate cheating but also because receivers might unknowingly smudge those parts of the response that most resemble the target, for instance as they trace a line with a finger or point out target-matching areas on the response sheet to another person. The judge's attention might be drawn to such areas because they look somewhat different to the rest of the response, not realising that the differences are due to handling or that the response sheet has been handled non-blind.

It is difficult to determine whether Experiment One employed sufficient measures to guard against this problem from the information presented in Lantz et al. (1994). They state that at the end of each trial (i) the receivers faxed a copy of their responses to Lantz, (ii) Lantz then sent the target, as feedback, to the receivers by return mail and (iii) the receivers sent their original responses to SAIC. However, Lantz et al. did not make it clear whether the receivers sent their responses to SAIC before or after they had received feedback from Lantz and whether the judge used the faxed or mailed response (a point also raised by Utts, 1995a).

RW asked May to clarify this point. He replied that it was possible that the receivers may have been in possession of their response and the feedback target at the same time (personal correspondence, 29 April 1996). May also noted that he had carried out the judging using a photocopy of the original version, not the faxed version (personal correspondence, 29 April

1996). He would therefore have been judging photocopies of some responses that may have been handled by the receivers after they had obtained feedback. However, May noted that Luke had compared the responses that the receivers faxed to Lantz with the originals that they sent to SAIC and found there were no differences between them (personal correspondence, 29 April 1996).

May later noted that this memory of events was incorrect, and that he had since discovered that he had judged from the faxes sent to Lantz, not the original responses (personal correspondence, 30 April 1997).

To summarise, Lantz et al. (1994) did not make it clear whether the judge used the faxed or original response. May first thought that he had used the originals, and that these were later compared with the faxed versions to ensure that the receivers had not altered their responses after obtaining the feedback target. However, May later stated that this was incorrect and that he actually judged from the faxes sent to Lantz. If this final reconstruction of events is correct, it is not clear why the experimenters compared the responses faxed to Lantz with the original response sent directly to SAIC. However, assuming that it is correct, the possibility of receivers inadvertently leaving any cues on their responses is ruled out. It does, however, open up a new potential pathway for sensory leakage, namely that Lantz (who knew those targets in the 'sender' condition) may have left cues on the faxes when he handled them directly after the completion of the trial. If this were the case, Lantz would only have left such cues on responses in the 'sender' condition, as he would be unaware of the targets in the 'no-sender' condition. One would therefore predict that the 'sender' trials would outscore the 'no sender' trials, but in fact there was no significant difference in scoring between the two sets of trials.

(iii) Safeguards against sensory leakage between the experimenter(s) and judge

As previously mentioned, in ESP research it is generally recognised that anyone who knows the identity of the target should not transcribe or edit the receivers' responses before they are given to a judge. In addition, some experimenters also describe precautions against non-

blind experimenters having any contact with the judge. This is to minimise the possibility of the judge obtaining information about the target without realising it through nonverbal cues, inadvertent references to target content etc. (Milton & Wiseman, 1997).

Lantz et al. (1994) did not document how much contact the experimenter(s) who knew the identities of the targets had with the judge (May). However, May (personal communication, 22 February 1997) informed RW that the SAIC personnel involved in the project (i.e., Luke, DeGraff and Macgowan) were working in offices a few meters away from his office and often had contact with him. In addition, although May had apparently told Utts that he had had no contact with Lantz (Utts, 1995a), he informed RW that the two of them occasionally met up (approximately three times a year) when Lantz visited SAIC (personal communication, 30 April 1997). However, May also stressed that the experimenters never discussed ongoing experiments with involved with the project (personal communication, 29 April 1996).

It is difficult to assess whether this agreement not to discuss the study with the judge would have been fully effective as a safeguard. Psychologists know little about the full range of ways in which inadvertent cueing can affect the behaviour of participants in experiments but there is no doubt that it happens (Rosenthal & Rubin, 1978). In the present case, an experimenter who knew the target identities might know that a particular image was the target a relatively large number of times and become biased towards referring to target-relevant content in his or her everyday conversation with the judge without realising it. Alternatively, the judge may have unknowingly overheard the experimenters discussing, for example, that a particular image or videotape was the target in several trials.

The possible impact of this potential problem is extremely difficult to gauge. Clearly, no accidental, blatant references to target material (such as 'by the way, one of the targets was Niagara Falls') could have been made by an experimenter, or May and the other experimenters would have realised that the protocol had been breached and would have taken appropriate steps. Whether or not there were subtle cues of the type discussed, undetectable

by the experimenters, is impossible to assess. The evaluation of this problem is further complicated by the difficulties discussed in the first section involved in establishing which experimenters knew the targets' identities and had contact with the judge, either directly or via intermediaries.

(iv) Safeguards against cues to past targets from responses judged out of their original order
Problems (discussed below) can arise in open-deck studies with trial-by-trial feedback of the target's identity, such as Experiment One, if the trials are judged out of order. Lantz et al. (1994) did not document whether the trials were judged in the order in which they were carried out, so RW asked May for more information. May (personal communication to RW, 29 April 1996) stated that trials were judged in random order rather than in their original order, about ten at any one time.

Kennedy (1979) has noted that such a procedure could introduce an artifact. His argument is best illustrated by the following scenario. Imagine a small experiment of only three trials. In the first trial the receiver draws a picture and then obtains feedback and discovers that the target was a picture of a waterfall. On trial two the receiver can't get the waterfall out of his or her head and draws an exact copy of it. The receiver then finds out that the target for this second trial was a picture of a mountain. In trial three the receiver draws an exact copy of this mountain and finds out that the target was a picture of a camel. If these trials were judged in the order in which they were carried out the receiver would be likely to score only at chance.

However, imagine that the trials were presented to the judge in reverse order. The judge would see the third trial first. He would see the receiver's drawing of the mountain and probably not assign the actual target (the camel) a first place rank. However, when the judge examined trial two he would look at the five possible targets and instantly recognise the mountain from the receiver's previous drawing. This might cause him to ignore the receiver's response for this trial (the waterfall) and choose the mountain. This trial would become a hit. When judging the first trial the judge would examine the five possible targets, immediately

recognise the waterfall from the receiver's drawing in the previously judged trials and choose the correct target. Again, this would become a hit.

In short, if receivers include elements of previous targets in their drawings on later trials, and the trials are judged out of their original order, the judge could pick up cues to the targets used in earlier trials from responses made in later trials. In this fictional example, the cues would have been so glaring that the judge would have seen what was happening and would have alerted the other experimenters to the problem. However, in a real, rather than fictional experiment the cues are likely to be less glaring and less likely to alert a judge to their nature. Thus, May could have been exposed to such cues and have been influenced by them without realising what they were. Although the trials that May judged were in random, rather than exact reverse order, quite a few trials in the study will have been judged before trials that actually took place earlier. Because of this problem, Kennedy (1979) recommended that whenever receivers are given trial-by-trial feedback, trials should be judged in the order in which they were carried out.

Discussion

To properly assess any experiment it is important to be able to accurately reconstruct its protocol. Unfortunately, the written sources describing Experiment One (Lantz et al, 1994; May and Lantz, 1991) do not contain information about several key aspects of the experiment's procedure. For example, they do not describe whether anyone who knew the identity of the target came into contact with the receivers' responses, how much contact the experimenters had with the judge, whether May judged from the original responses or the faxes and whether the trials were judged in order. The Hyman/Honorton joint communiqué recommended that research analysts should be able to reconstruct experimental procedures from the descriptions provided in written reports (Hyman & Honorton, 1986, p.360). Future research should concentrate on developing strategies which would help parapsychologists provide a more complete, unambiguous and reliable description of their studies.

Second, as the written reports did not provide very much information about certain aspects of the experiment, it was necessary to try to reconstruct the protocol on the basis of the experimenters' memories. The information provided by these individuals proved extremely valuable and the authors are indebted to the experimenters for taking the time to make these details available. However, the experiment took place four years ago and was only one of six remote viewing experiments conducted by the research team. As such, it is perhaps not surprising that these individuals had problems remembering all of the details needed to fully reconstruct the experiment. For example, May presented several versions of how many experimenters were involved in the study and their roles at SAIC. May also first thought that he judged from the original responses, only to later report that he had actually used the faxes. This further emphasises the need for investigators to make accurate records of studies when they happen, rather than having to depend upon their own memories (which are more likely to be influenced by the effects of bias and difficulties of recall) at a later date.

Third, when the reporting deficiencies discussed above were resolved, two potential pathways for sensory leakage remained. The first concerned the possibility of sensory leakage from experimenters to judge, the second arose through the trials being judged out of order. Such possible sources of leakage are not trivial - the outcome of the experiment as a whole, and even of its most successful subgroup of trials, is not statistically sufficiently robust that any of these possible flaws appear unlikely as explanations of the results.

Fourth, all of the opportunities for sensory cueing that we have identified involve some form of information leakage from non-blind personnel to the judge. This raises the question of whether rejudging under blind conditions could answer the question of whether the flaws we have described did, in fact, contribute to the study's outcome. Such a reanalysis may be possible, but would have to be very carefully thought out. One possibility, suggested by two anonymous referees of this paper, would involve having someone who is blind to the targets copy the viewers drawings (not including any smudges or other possible handling cues) onto new response sheets, and then having these blind judged in the original order .

Fifth, the question arises as to why this experiment was poorly recorded and incorporated design flaws that had been discussed in the literature before the SAIC research began (Milton & Wiseman, 1997). This is not an easy question to answer. The research program was a relatively costly one. It was not the first set of remote viewing studies conducted by this laboratory - indeed, it built upon a previous set of studies that had been criticised on grounds of sensory leakage [see Footnote 1]. Also, the experiment had been endorsed by a 'Scientific Oversight Committee' prior to being carried out. All of these factors should have led to a properly conducted and well written up experiment.

Sixth, the AIR assessors state that they often had to rely upon the May's memory for details concerning important aspects of the SAIC experiments (Utts, 1995a). This is worrying, given that May's initial answers to many of RW's questions were incorrect. For example, May inaccurately recalled the role of the SAIC personal in the experiment and whether he judged from the original responses or the faxes. This problem is reflected in the fact that many of the procedural details presented in the AIR report are, according to May's latest recall, incorrect. For example, Utts stated that May had no contact with the Lantz and implied that May judged the original responses sent to SAIC. May now claims that neither are true. In addition, this critique has demonstrated that the AIR assessors missed two flaws in the experimental design.

It is important to place this re-analysis in perspective. Compared to many other published free-response ESP studies, Experiment One is not badly conducted. It is not uncommon, for example, for independent judges to judge receivers' responses out of their original order or for authors not to describe any precautions against contact between non-blind experimenters and judges. Indeed, Experiment One includes safeguards by no means always reported by other studies. For example, the statistical analyses to be carried out were described as having been prespecified in advance of the experiment and the large distances between sender and receiver make it unlikely that there would have been any sensory leakage between them during the experimental session (something not always clear in other experimental reports when researchers attempt to convey details of the layout and construction of their

laboratories that should rule out sensory cues). Experiment One has only attracted the level of detailed critique it has received here because of the strong and very public claims made for its success and methodological rigour, not because it appeared particularly weak.

Although we may differ in our evaluations of the SAIC work, nevertheless the SAIC experimenters and the present authors all share the same goal of wanting to determine whether it is possible to show empirical evidence for a remote viewing effect under conditions that rule out information leakage. Through their co-operation, the SAIC team have contributed to the process of refining remote viewing methodology so that researchers who attempt to replicate their work can avoid the problems that they encountered. If their results do in fact reflect a genuine remote viewing effect, they have also given those researchers a head start in achieving success in their studies, by means of their process-oriented approach. Just as a collaborative critique of ESP ganzfeld studies (Hyman & Honorton, 1986) helped to raise standards of experimental conduct and reporting in that area of research, we hope that the joint effort reported here will have a similar effect for the future remote viewing studies that are bound to follow.

References

- Akers, C. (1984). Methodological criticisms of parapsychology. In S. Krippner (Ed.), *Advances in parapsychological research 4* (pp.112-164). New York: Plenum Press.
- Hyman, R. (1995a). Evaluation of the program on anomalous mental phenomena, In M.D. Mumford, A.M. Rose and D.A. Goslin. *An evaluation of remote viewing: Research and Applications* (pp 3.43-3.78). Prepared by The American Institutes for Research.
- Hyman, R. (1995b). Evaluation of the program on anomalous mental phenomena. *Journal of Parapsychology*, 59(4), 321-351.

Hyman, R. (1996). Evaluation of the program on anomalous mental phenomena. *Journal of Scientific Exploration*, 10(1), 31-58.

Hyman, R. and Honorton, C. (1986). A joint communiqué: The psi ganzfeld controversy. *Journal of Parapsychology*, 50, 350-364.

Kennedy, J.E. (1979). More on methodological issues in free-response psi experiments. *Journal of the American Society for Psychical Research*, 73, 395-401.

Lantz, N.D., Luke, W.L. and May, E.C. (1994). Target and sender dependencies in anomalous cognition experiments. *Journal of Parapsychology*, 58, 285-302.

Marks, D. and Kammann, R. (1980). *The Psychology of the Psychic*. Buffa, USA: Prometheus Books.

Marks, D. (1981a). On the review of *The Psychology of the Psychic*: a reply to Dr Morris. *Journal of the American Society for Psychical Research*, 75, 197-203.

Marks, D. (1981b). Sensory cues invalidate remote viewing experiments. *Nature*, 292, 177.

May, E.C. and Lantz, N.D. (1991). *Target and sender dependencies in anomalous cognition: Report submitted to The Scientific Oversight Committee*. Science Applications International Corporation.

Milton, J. (1993). A meta-analysis of waking state of consciousness, free response ESP studies. In *The Parapsychological Association 36th Annual Convention: Proceedings of presented papers*, pp. 87-104.

Milton, J. (1995). Issues in the adoption of methodological safeguards in parapsychology experiments. In *The Parapsychological Association 38th Annual Convention: Proceedings of presented papers*. ppxx-xx.

Milton, J. and Wiseman, R.J. (1997). *Guidelines for research into extrasensory perception*. Hatfield, UK: University of Hertfordshire Press.

Morris, R.L. (1986). What psi is not: The necessity for experiments. In H.L. Edge, R.L. Morris, J. Palmer and J.H. Rush (Eds.), *Foundations of parapsychology: Exploring the boundaries of human capability* (pp.70-110). Boston: Routledge and Kegan Paul.

Morris, R.L. (1980). Some comments on the assessment of parapsychological studies: a review of *The Psychology of the Psychic* by D. Marks and R. Kammann, *Journal of the American Society for Psychical Research*, 74, 425-43.

Mumford, M.D., Rose, A.M. and Goslin, D.A. (1995). *An evaluation of remote viewing: Research and Applications*. Prepared by The American Institutes for Research.

Puthoff, H. and Targ, R. (1981). Rebuttal of criticisms of remote viewing experiments, *Nature*, 292, 388.

Rosenthal, R. and Rubin, D.B. (1978). Interpersonal expectancy effects: The first 345 studies. *The Behavioral and Brain Sciences*, 3, 377-386.

Solfvin, G.F., Kelly, E.F., & Burdick, D.S. (1978). Some new methods of analysis for preferential-ranking data. *Journal of the Society for Psychical Research*, 72, 93-109.

Targ, R. and Puthoff, H. (1974). Information transmission under conditions of sensory shielding. *Nature*, 252, 602-607.

Targ, R. and Puthoff, H. (1977). *Mind-reach*. New York: Delacorte.

Stanford, R.G. and Stein, A.G. (1994). A meta-analysis of ESP studies contrasting hypnosis and a comparison condition. *Journal of Parapsychology*, 58, 235-269.

Utts, J. (1995a). An assessment of the evidence for psychic functioning. *Journal of Parapsychology*, In M.D. Mumford, A.M. Rose and D.A. Goslin. *An evaluation of remote viewing: Research and Applications* (pp 3.2-3.42). Prepared by The American Institutes for Research.

Utts, J. (1995b). An assessment of the evidence for psychic functioning. *Journal of Parapsychology*, 59(4), 289-320.

Utts, J. (1996). An assessment of the evidence for psychic functioning. *Journal of Scientific Exploration*, 10(1), 3-30.

Footnotes

1. For the debate about these studies see Targ and Puthoff (1974, 1977), Marks and Kammann (1980), Marks (1981a, b), Puthoff and Targ (1981) and Morris (1980).
2. Lantz et al. (1994) do not apply the continuity correction recommended by Solfvin et al., and our probability estimates are therefore slightly more conservative than those that they report.
3. All personal communications cited in this paper consist of emails. Although Dr May did not object to his emails being cited verbatim in an earlier version of this paper (presented at the 1997 Convention of the Parapsychological Association), he has asked the authors to paraphrase them here.