

An Assessment of the Evidence for Psychic Functioning

JESSICA UTTS

Division of Statistics, University of California, Davis, CA 95616

Abstract — Research on psychic functioning, conducted over a two decade period, is examined to determine whether or not the phenomenon has been scientifically established. A secondary question is whether or not it is useful for government purposes. The primary work examined in this report was government sponsored research conducted at Stanford Research Institute, later known as SRI International, and at Science Applications International Corporation, known as SAIC.

Using the standards applied to any other area of science, it is concluded that psychic functioning has been well established. The statistical results of the studies examined are far beyond what is expected by chance. Arguments that these results could be due to methodological flaws in the experiments are soundly refuted. Effects of similar magnitude to those found in government-sponsored research at SRI and SAIC have been replicated at a number of laboratories across the world. Such consistency cannot be readily explained by claims of flaws or fraud.

The magnitude of psychic functioning exhibited appears to be in the range between what social scientists call a small and medium effect. That means that it is reliable enough to be replicated in properly conducted experiments, with sufficient trials to achieve the long-run statistical results needed for replicability.

A number of other patterns have been found, suggestive of how to conduct more productive experiments and applied psychic functioning. For instance, it doesn't appear that a sender is needed. Precognition, in which the answer is known to no one until a future time, appears to work quite well. Recent experiments suggest that if there is a psychic sense then it works much like our other five senses, by detecting change. Given that physicists are currently grappling with an understanding of time, it may be that a psychic sense exists that scans the future for major change, much as our eyes scan the environment for visual change or our ears allow us to respond to sudden changes in sound.

It is recommended that future experiments focus on understanding how this phenomenon works, and on how to make it as useful as possible. There is little benefit to continuing experiments designed to offer proof, since there is little more to be offered to anyone who does not accept the current collection of data.

1. Introduction

This paper was written for the American Institutes of Research, which had been commissioned by the CIA, at the request of Congress, to evaluate the government program in remote viewing and related areas. We were asked to limit our investigation to the government work. Therefore, this report is not intended to be an overall assessment of the field of parapsychology, or even of all of the work in remote viewing or anomalous cognition. A report written for a more general audience, rather than specifically requested by the CIA and Congress, would have rightfully had a broader focus than the one we were directed to use here. My conclusions in this report are substantiated by the limited review provided, but are firmly supported by broader reviews previously published by myself and others.

The purpose of this report is to examine a body of evidence collected over the past few decades in an attempt to determine whether or not psychic functioning is possible. Secondary questions include whether or not such functioning can be used productively for government purposes, and whether or not the research to date provides any explanation for how it works.

There is no reason to treat this area differently from any other area of science that relies on statistical methods. Any discussion based on belief should be limited to questions that are not data-driven, such as whether or not there are any methodological problems that could substantially alter the results. It is too often the case that people on both sides of the question debate the existence of psychic functioning on the basis of their personal belief systems rather than on an examination of the scientific data.

One objective of this report is to provide a brief overview of recent data as well as the scientific tools necessary for a careful reader to reach his or her own conclusions based on that data. The tools consist of a rudimentary overview of how statistical evidence is typically evaluated, and a listing of methodological concerns particular to experiments of this type.

Government-sponsored research in psychic functioning dates back to the early 1970s when a program was initiated at what was then the Stanford Research Institute, now called SRI International. That program was in existence until 1989. The following year, government sponsorship moved to a program at Science Applications International Corporation (SAIC) under the direction of Dr. Edwin May, who had been employed in the SRI program since the mid 1970s and had been Project Director from 1986 until the close of the program.

This report will focus most closely on the most recent work, done by SAIC. Section 2 describes the basic statistical and methodological issues required to understand this work; Section 3 discusses the program at SRI; Section 4 covers the SAIC work (with some of the details in an Appendix); Section 5 is concerned with external validation by exploring related results from other laboratories; Section 6 includes a discussion of the usefulness of this capability for government purposes and Section 7 provides conclusions and recommendations.

2. Science Notes

2.1 Definitions and Research Procedures

There are two basic types of functioning that are generally considered under the broad heading of psychic or paranormal abilities. These are classically known as extrasensory perception (ESP), in which one acquires information through unexplainable means and psychokinesis, in which one physically manipulates the environment through unknown means. The SAIC laboratory uses more neutral terminology for these abilities; they refer to ESP as anomalous cognition (AC) and to psychokinesis as anomalous perturbation (AP). The vast majority of work at both SRI and SAIC investigated anomalous cognition rather than anomalous perturbation, although there was some work done on the latter.

Anomalous cognition is further divided into categories based on the apparent source of the information. If it appears to come from another person, the ability is called telepathy, if it appears to come in real time but not from another person it is called clairvoyance and if the information could have only been obtained by knowledge of the future, it is called precognition.

It is possible to identify apparent precognition by asking someone to describe something for which the correct answer isn't known until later in time. It is more difficult to rule out precognition in experiments attempting to test telepathy or clairvoyance, since it is almost impossible to be sure that subjects in such experiments never see the correct answer at some point in the future. These distinctions are important in the quest to identify an explanation for anomalous cognition, but do not bear on the existence issue.

The vast majority of anomalous cognition experiments at both SRI and SAIC used a technique known as remote viewing. In these experiments, a viewer attempts to draw or describe (or both) a target location, photograph, object or short video segment. All known channels for receiving the information are blocked. Sometimes the viewer is assisted by a monitor who asks the viewer questions; of course in such cases the monitor is blind to the answer as well. Sometimes a sender is looking at the target during the session, but sometimes there is no sender. In most cases the viewer eventually receives feedback in which he or she learns the correct answer, thus making it difficult to rule out precognition as the explanation for positive results, whether or not there was a sender.

Most anomalous cognition experiments at SRI and SAIC were of the *free-response* type, in which viewers were simply asked to describe the target. In contrast, a *forced-choice* experiment is one in which there are a small number of known choices from which the viewer must choose. The latter may be easier to evaluate statistically but they have been traditionally less successful than free-response experiments. Some of the work done at SAIC addresses potential explanations for why that might be the case.

2.2 Statistical Issues and Definitions

Few human capabilities are perfectly replicable on demand. For example, even the best hitters in the major baseball leagues cannot hit on demand. Nor can we predict when someone will hit or when they will score a home run. In fact, we cannot even predict whether or not a home run will occur in a particular game. That does not mean that home runs don't exist.

Scientific evidence in the statistical realm is based on replication of the same average performance or relationship *over the long run*. We would not expect a fair coin to result in five heads and five tails over each set of ten tosses, but we can expect the proportion of heads and tails to settle down to about one half over a very long series of tosses. Similarly, a good baseball hitter will not hit the ball exactly the same proportion of times in each game but should be relatively consistent over the long run.

The same should be true of psychic functioning. Even if there truly is an effect, it may never be replicable on demand in the short run even if we understand how it works. However, over the long run in well-controlled laboratory experiments we should see a consistent level of functioning, above that expected by chance. The anticipated level of functioning may vary based on the individual players and the conditions, just as it does in baseball, but given players of similar ability tested under similar conditions the results should be replicable over the long run. In this report we will show that replicability in that sense has been achieved.

2.2.1 P-values and Comparison with Chance. In any area of science, evidence based on statistics comes from comparing what actually happened to what should have happened by chance. For instance, without any special interventions about 51 percent of births in the United States result in boys. Suppose someone claimed to have a method that enabled one to increase the chances of having a baby of the desired sex. We could study their method by comparing how often births resulted in a boy when that was the intended outcome. If that percentage was higher than the chance percentage of 51 percent *over the long run*, then the claim would have been supported by statistical evidence.

Statisticians have developed numerical methods for comparing results to what is expected by chance. Upon observing the results of an experiment, the *p-value* is the answer to the following question: *If chance alone is responsible for the results, how likely would we be to observe results this strong or stronger?* If the answer to that question, i.e. the p-value is very small, then most researchers are willing to rule out chance as an explanation. In fact it is commonly accepted practice to say that if the p-value is 5 percent (*0.05*) or less, then we can rule out chance as an explanation. In such cases, the results are said to be *statistically significant*. Obviously the smaller the p-value, the more convincingly chance can be ruled out.

Notice that when chance alone is at work, we *erroneously* find a statistically significant result about 5 percent of the time. For this reason and others, most

reasonable scientists require replication of non-chance results before they are convinced that chance can be ruled out.

2.2.2 Replication and Effect Sizes. In the past few decades scientists have realized that true replication of experimental results should focus on the magnitude of the effect, or the *effect* size rather than on replication of the p-value. This is because the latter is heavily dependent on the size of the study. In a very large study, it will take only a small magnitude effect to convincingly rule out chance. In a very small study, it would take a huge effect to convincingly rule out chance.

In our hypothetical sex-determination experiment, suppose 70 out of 100 births designed to be boys actually resulted in boys, for a rate of 70 percent instead of the 51 percent expected by chance. The experiment would have a p-value of 0.0001, quite convincingly ruling out chance. Now suppose someone attempted to replicate the experiment with only ten births and found 7 boys, i.e. also 70 percent. The smaller experiment would have a p-value of 0.19, and would not be statistically significant. If we were simply to focus on that issue, the result would appear to be a failure to replicate the original result, even though it achieved exactly the same 70 percent boys! In only ten births it would require 90 percent of them to be boys before chance could be ruled out. Yet the 70 percent rate is a more exact replication of the result than the 90 percent.

Therefore, while p-values should be used to assess the overall evidence for a phenomenon, they should not be used to define whether or not a replication of an experimental result was "successful." Instead, a successful replication should be one that achieves an effect that is within expected statistical variability of the original result, or that achieves an even stronger effect for explainable reasons.

A number of different *effect* size measures are in use in the social sciences, but in this report we will focus on the one used most often in remote viewing at SRI and SAIC. Because the definition is somewhat technical it is given in Appendix 1. An intuitive explanation will be given in the next subsection. Here, we note that an effect size of 0 is consistent with chance, and social scientists have, by convention, declared an effect size of 0.2 as small, 0.5 as medium and 0.8 as large. A medium effect size is supposed to be visible to the naked eye of a careful observer, while a large effect size is supposed to be evident to any observer.

2.2.3 Randomness and Rank-Order Judging. At the heart of any statistical method is a definition of what should happen "randomly" or "by chance." Without a random mechanism, there can be no statistical evaluation.

There is nothing random about the responses generated in anomalous cognition experiments; in other words, there is no way to define what they would look like "by chance." Therefore, the random mechanism in these experiments must be in the choice of the target. In that way, we can compare the response to the target and answer the question: "If chance alone is at work, what is the

probability that a *target* would be chosen that matches this *response* as well as or better than does the actual target?"

In order to accomplish this purpose, a properly conducted experiment uses a set of targets defined in advance. The target for each remote viewing is then selected randomly, in such a way that the probability of getting each possible target is known.

The SAIC remote viewing experiments and all but the early ones at SRI used a statistical evaluation method known as *rank-order judging*. After the completion of a remote viewing, a judge who is blind to the true target (called a *blind judge*) is shown the response and five potential targets, one of which is the correct answer and the other four of which are "decoys." Before the experiment is conducted each of those five choices must have had an equal chance of being selected as the actual target. The judge is asked to assign a rank to each of the possible targets, where a rank of one means it matches the response most closely, and a rank of five means it matches the least.

The rank of the correct target is the numerical score for that remote viewing. By chance alone the actual target would receive each of the five ranks with equal likelihood, since despite what the response said the target matching it best would have the same chance of selection as the one matching it second best and so on. The average rank by chance would be three. Evidence for anomalous cognition occurs when the average rank over a series of trials is significantly lower than three. (Notice that a rank of one is the best possible score for each viewing.)

This scoring method is conservative in the sense that it gives no extra credit for an excellent match. A response that describes the target almost perfectly will achieve the same rank of one as a response that contains only enough information to pick the target as the best choice out of the five possible choices. One advantage of this method is that it is still valid even if the viewer knows the set of possible targets. The probability of a first place match by chance would still be only one in five. This is important because the later SRI and many of the SAIC experiments used the same large set of *National Geographic* photographs as targets. Therefore, the experienced viewers would eventually become familiar with the range of possibilities since they were usually shown the answer at the end of each remote viewing session.

For technical reasons explained in Appendix 1, the effect size for a series of remote viewings using rank-order judging with five choices is $(3.0 - \text{average rank})/2^{1/2}$. Therefore, small, medium and large effect sizes (0.2, 0.5 and 0.8) correspond to average ranks of 2.72, 2.29 and 1.87, respectively. Notice that the largest effect size possible using this method is 1.4, which would result if every remote viewing achieved a first place ranking.

2.3 Methodological Issues

One of the challenges in designing a good experiment in any area of science

is to close the loopholes that would allow explanations other than the intended one to account for the results.

There are a number of places in remote viewing experiment where information could be conveyed by normal means if proper precautions are not taken. The early SRI experiments suffered from some of those problems, but the later SRI experiments and the SAIC work were done with reasonable methodological rigor, with some exceptions noted in the detailed descriptions of the SAIC experiments in Appendix 2.

The following list of methodological issues shows the variety of concerns that must be addressed. It should be obvious that a well-designed experiment requires careful thought and planning:

- No one who has knowledge of the specific target should have any contact with the viewer until after the response has been safely secured.
- No one who has knowledge of the specific target or even of whether or not the session was successful should have any contact with the judge until after that task has been completed.
- No one who has knowledge of the specific target should have access to the response until after the judging has been completed.
- Targets and decoys used in judging should be selected using a well-tested randomization device.
- Duplicate sets of targets photographs should be used, one during the experiment and one during the judging, so that no cues (like fingerprints) can be inserted onto the target that would help the judge recognize it.

The criterion for stopping an experiment should be defined in advance so that it is not called to a halt when the results just happen to be favorable. Generally, that means specifying the number of trials in advance, but some statistical procedures require or allow other stopping rules. The important point is that the rule be defined in advance in such a way that there is no ambiguity about when to stop.

- Reasons, if any, for excluding data must be defined in advance and followed consistently, and should not be dependent on the data. For example, a rule specifying that a trial could be aborted if the viewer felt ill would be legitimate, but only if the trial was aborted before anyone involved in that decision knew the correct target.
- Statistical analyses to be used must be planned in advance of collecting the data so that a method most favorable to the data isn't selected *post hoc*. If multiple methods of analysis are used the corresponding conclusions must recognize that fact.

2.4 *Prima Facie Evidence*

According to *Webster's Dictionary*, in law *prima facie* evidence is "evidence having such a degree of probability that it must prevail unless the contrary be proved." There are a few examples of applied, non-laboratory remote viewings

provided to the review team that would seem to meet that criterion for evidence. These are examples in which the sponsor or another government client asked for a single remote viewing of a site, known to the requester in real time or in the future, and the viewer provided details far beyond what could be taken as a reasonable guess. Two such examples are given by May (1995) in which it appears that the results were so striking that they far exceed the phenomenon as observed in the laboratory. Using a *post hoc* analysis, Dr. May concluded that in one of the cases the remote viewer was able to describe a microwave generator with 80 percent accuracy, and that of what he said almost 70 percent of it was reliable. Laboratory remote viewings rarely show that level of correspondence.

Notice that standard statistical methods cannot be used in these cases because there is no standard for probabilistic comparison. But evidence gained from applied remote viewing cannot be dismissed as inconsequential just because we cannot assign specific probabilities to the results. It is most important to ascertain whether or not the information was achievable in other standard ways. In Section 3 an example is given in which a remote viewer allegedly gave code words from a secret facility that he should not have even known existed. Suppose the sponsors could be absolutely certain that the viewer could not have known about those code words through normal means. Then even if we can't assign an exact probability to the fact that he guessed them correctly, we can agree that it would be very small. That would seem to constitute *prima facie* evidence unless an alternative explanation could be found. Similarly, the viewer who described the microwave generator allegedly knew only that the target was a technical site in the United States. Yet, he drew and described the microwave generator, including its function, its approximate size, how it was housed and that it had "a beam divergence angle of 30 degrees" (May, 1995, p. 15).

Anecdotal reports of psychic functioning suffer from a similar problem in terms of their usefulness as proof. They have the additional difficulty that the "response" isn't even well-defined in advance, unlike in applied remote viewing where the viewer provides a fixed set of information on request. For instance, if a few people each night happen to dream of plane crashes, then some will obviously do so on the night before a major plane crash. Those individuals may interpret the coincidental timing as meaningful. This is undoubtedly the reason many people think the reality of psychic functioning is a matter of belief rather than science, since they are more familiar with the provocative anecdotes than with the laboratory evidence.

3. The SRI Era

3.1 Early Operational Successes and Evaluation

According to Puthoff and Targ (1975) the scientific research endeavor at SRI may never have been supported had it not been for three apparent opera-

tional successes in the early days of the program. These are detailed by Puthoff and Targ (1975), although the level of the matches is not clearly delineated.

One of the apparent successes concerned the "West Virginia Site" in which two remote viewers purportedly identified an underground secret facility. One of them apparently named codewords and personnel in this facility accurately enough that it set off a security investigation to determine how that information could have been leaked. Based only on the coordinates of the site, the viewer first described the above ground terrain, then proceeded to describe details of the hidden underground site.

The same viewer then claimed that he could describe a similar Communist Bloc site and proceeded to do so for a site in the Urals. According to Puthoff and Targ "the two reports for the West Virginia Site, and the report for the Urals Site were verified by personnel in the sponsor organization as being substantially correct (p. 8)."

The third reported operational success concerned an accurate description of a large crane and other information at a site in Semipalatinsk, USSR. Again the viewer was provided with only the geographic coordinates of the site and was asked to describe what was there.

Although some of the information in these examples was verified to be highly accurate, the evaluation of operational work remains difficult, in part because there is no chance baseline for comparison (as there is in controlled experiments) and in part because of differing expectations of different evaluators. For example, a government official who reviewed the Semipalatinsk work concluded that there was no way the remote viewer could have drawn the large gantry crane unless "he actually saw it through remote viewing, or he was informed of what to draw by someone knowledgeable of [the site]." Yet that same analyst concluded that "the remote viewing of [the site] by subject S1 proved to be unsuccessful" because "the only positive evidence of the rail-mounted gantry crane was far outweighed by the large amount of negative evidence noted in the body of this analysis." In other words, the analyst had the expectation that in order to be "successful" a remote viewing should contain accurate information only.

Another problem with evaluating this operational work is that there is no way to know with certainty that the subject did not speak with someone who had knowledge of the site, however unlikely that possibility may appear. Finally, we do not know to what degree the results in the reports were selectively chosen because they were correct. These problems can all be avoided with well designed controlled experiments.

3.2 The Early Scientific Effort at SRI

During 1974 and early 1975 a number of controlled experiments were conducted to see if various types of target material could be successfully described with remote viewing. The results reported by Puthoff and Targ (1975) indicated success with a wide range of material, from "technical" targets like a xerox

machine to natural settings, like a swimming pool. But these and some of the subsequent experiments were criticized on statistical and methodological grounds; we briefly describe one of the experiments and criticisms of it to show the kinds of problems that existed in the early scientific effort.

The largest series during the 1973 to 1975 time period involved remote viewing of natural sites. Sites were randomly selected for each trial from a set of 100 possibilities. They were selected "without replacement," meaning that sites were not reused once they had been selected. The series included eight viewers, including two supplied by the sponsor. Many of the descriptions showed a high degree of subjective correspondence, and the overall statistical results were quite striking for most of the viewers.

Critics attacked these experiments on a number of issues, including the selection of sites without replacement and the statistical scoring method used. The results were scored by having a blind judge attempt to match the target material with the transcripts of the responses. A large fraction of the matches were successful. But critics noted that some successful matching could be attained just from cues contained in the transcripts of the material, like when a subject mentioned in one session what the target had been in the previous session. Because sites were selected without replacement, knowing what the answer was on one day would exclude that target site from being the answer on any other day. There was no way to determine the extent to which these problems influenced the results. The criticisms of these and subsequent experiments, while perhaps unwelcome at the time, have resulted in substantially improved methodology in these experiments.

3.3 An Overall Analysis of the SRI Experiments: 1973-1988

In 1988 an analysis was made of all of the experiments conducted at SRI from 1973 until that time (May et al, 1988). The analysis was based on all 154 experiments conducted during that era, consisting of over 26,000 individual trials. Of those, almost 20,000 were of the forced choice type and just over a thousand were laboratory remote viewings. There were a total of 227 subjects in all experiments.

The statistical results were so overwhelming that results that extreme or more so would occur only about once in every 10^{20} such instances if chance alone is the explanation (i.e., the p -value was less than 10^{-20}). Obviously some explanation other than chance must be found. Psychic functioning may not be the only possibility, especially since some of the earlier work contained methodological problems. However, the fact that the same level of functioning continued to hold in the later experiments, which did not contain those flaws, lends support to the idea that the methodological problems cannot account for the results. In fact, there was a talented group of subjects (labeled G1 in that report) for whom the effects were stronger than for the group at large. According to Dr. May, the majority of experiments with that group were conducted later in the program, when the methodology had been substantially improved.

In addition to the statistical results, a number of other questions and patterns were examined. A summary of the results revealed the following:

1. "Free response" remote viewing, in which subjects describe a target, was much more successful than "forced choice" experiments, in which subjects were asked to choose from a small set of possibilities.
2. There was a group of six selected individuals whose performance far exceeded that of unselected subjects. The fact that these same selected individuals consistently performed better than others under a variety of protocols provides a type of replicability that helps substantiate the validity of the results. If methodological problems were responsible for the results, they should not have affected this group differently from others.
3. Mass-screening efforts found that about one percent of those who volunteered to be tested were consistently successful at remote viewing. This indicates that remote viewing is an ability that differs across individuals, much like athletic ability or musical talent. (Results of mass screenings were not included in the formal analysis because the conditions were not well-controlled, but the subsequent data from subjects found during mass-screening were included.)
4. Neither practice nor a variety of training techniques consistently worked to improve remote viewing ability. It appears that it is easier to find than to train good remote viewers.
5. It is not clear whether or not feedback (showing the subject the right answer) is necessary, but it does appear to provide a psychological boost that may increase performance.
6. Distance between the target and the subject does not seem to impact the quality of the remote viewing.
7. Electromagnetic shielding does not appear to inhibit performance.
8. There is compelling evidence that precognition, in which the target is selected after the subject has given the description, is also successful.
9. There is no evidence to support anomalous perturbation (psychokinesis), i.e. physical interaction with the environment by psychic means.

3.4 Consistency with Other Laboratories in the Same Era

One of the hallmarks of a real phenomenon is that its magnitude is replicable by various researchers working under similar conditions. The results of the overall SRI analysis are consistent with results of similar experiments in other laboratories. For instance, an overview of forced choice precognition experiments (Honorton and Ferrari, 1989) found an average "effect size" per experimenter of 0.033, whereas all forced choice experiments at SRI resulted in a similar effect size of 0.052. The comparison is not ideal since the SRI forced choice experiments were not necessarily precognitive and they used different types of target material than the standard card-guessing experiments.

Methodologically sound remote viewing has not been undertaken at other

laboratories, but a similar regime called the ganzfeld (described in more detail in Section 5) has shown to be similarly successful. The largest collection of ganzfeld experiments was conducted from 1983 to 1989 at the Psychophysical Research Laboratories in Princeton, NJ. Those experiments were also reported by separating novices from experienced subjects. The overall effect size for novice remote viewing at SRI was 0.164, while the effect size for novices in the ganzfeld at PRL was a very similar 0.17. For experienced remote viewers at SRI the overall effect size was 0.385; for experienced viewers in the ganzfeld experiments it was 0.35. These consistent results across laboratories help refute the idea that the successful experiments at any one lab are the result of fraud, sloppy protocols or some methodological problem and also provide an indication of what can be expected in future experiments.

4. The SAIC Era

4.1 An Overview

The review team decided to focus more intensively on the experiments conducted at Science Applications International Corporation (SAIC), because they provide a manageable yet varied set to examine in detail. They were guided by a Scientific Oversight Committee consisting of experts in a variety of disciplines, including a winner of the Nobel Prize in Physics, internationally known professors of statistics, psychology, neuroscience and astronomy and a medical doctor who is a retired U.S. Army Major General. Further, we have access to the details for the full set of SAIC experiments, unlike for the set conducted at SRI. Whatever details may be missing from the written reports are obtainable from the principal investigator, Dr. Edwin May, to whom we have been given unlimited access.

In a memorandum dated July 25, 1995, Dr. Edwin May listed the set of experiments conducted by SAIC. There were ten experiments, all designed to answer questions about psychic functioning, raised by the work at SRI and other laboratories, rather than just to provide additional proof of its existence. Some of the experiments were of a similar format to the remote viewing experiments conducted at SRI and we can examine those to see whether or not they replicated the SRI results. We will also examine what new knowledge can be gained from the results of the SAIC work.

4.2 The Ten Experiments

Of the ten experiments done at SAIC, six of them involved remote viewing and four did not. Rather than list the details in the body of this report, Appendix 2 gives a brief description of the experiments. What follows is a discussion of the methodology and results for the experiments as a whole. Because of the fundamental differences between remote viewing and the other types of experiments, we discuss them separately.

In the memorandum of 25 July 1995, Dr. May provided the review team

with details of the ten experiments, including a short title, number of trials, effect size and overall p-value for each one. His list was in time sequence. It is reproduced in Table 1, using his numbering system, with the experiments categorized by type, then sequentially within type. The effect size estimates are based on a limited number of trials, so they are augmented with an interval to show the probable range of the true effect (e.g. 0.124 ± 0.071 indicates a range from 0.053 to 0.195). Remember that an effect size of 0 represents chance, while a positive effect size indicates positive results.

TABLE 1
SAIC Experiments Listed by Dr. Edwin May

Expr	Title	Trials	Effect Size	p-value
Remote Viewing Experiments				
1	Target dependencies	200	0.124 ± 0.071	0.040
4	AC with binary coding	40	-0.067 ± 0.158	0.664
5	AC lucid dreams, base	24	0.088 ± 0.204	0.333
6	AC lucid dreams, pilot	21	0.368 ± 0.218	0.046
9	ERD AC Behavior	70	0.303 ± 0.120	0.006
10	Entropy II	90	0.550 ± 0.105	9.1×10^{-8}
Other Experiments				
2	AC of binary targets	300	0.123 ± 0.058	0.017
3	MEG Replication	12,000s	MCE	MCE
7	Remote observation	48	0.361 ± 0.144	0.006
8	ERD EEG investigation	7,000s	MCE	MCE

4.3 Assessing the Remote Viewing Experiments by Homogeneous Sets of Sessions

While Table 1 provides an overall assessment of the results of each experiment, it does so at the expense of information about variability among viewers and types of targets. In terms of understanding the phenomenon, it is important to break the results down into units that are as homogeneous as possible in terms of procedure, individual viewer and type of target. This is also important in order to assess the impact of any potential methodological problems. For example, in one pilot experiment (E6, AC in Lucid Dreams) viewers were permitted to take the targets home with them in sealed envelopes. Table 2 presents the effect size results at the most homogeneous level possible based on the information provided. For descriptions of the experiments, refer to Appendix 2. Overall effect sizes for each viewer and total effect sizes for each experiment are weighted according to the number of trials, so each trial receives equal weight.

TABLE 2
Individual Effect Sizes

Experiment	Expert Remote Viewers						Total
	009	131	372	389	518	Unknown/Other	
	Static Targets (National Geographies)						
E1: Static	0.424	-0.071	0.424	0.177	0.283	NA	0.247
E9	0.432	NA	0.354	0.177	NA	NA	0.303
E10: Static	0.566	NA	0.801	-0.071	0.778	NA	0.550
E5 (see note)	NA	NA	NA	NA	NA	0.088	0.088
E6 (see note)	NA	NA	NA	NA	NA	0.370	0.370
E4 (see note)	-0.112	NA	0.000	0.112	NA	-0.559	-0.067
	Dynamic Targets (Video Film Clips)						
E1: Dynamic	0.000	0.354	-0.283	0.000	-0.071	NA	0.000
E10: Dynamic	0.919	NA	0.754	0.000	0.424	NA	0.550
Overall	0.352	0.141	0.340	0.090	0.271	NA	

Note: Experiment 5 did not include any expert viewers. Experiment 6 included 4 expert viewers but separate results were not provided. Experiment 4 used a specially designed target set and only 4 choices in judging.

4.4 Consistency and Replicability of the Remote Viewing Results

One of the most important hallmarks of science is replicability. A phenomenon with statistical variability, whether it is scoring home runs in baseball, curing a disease with chemotherapy or observing psychic functioning, should exhibit about the same level of success in the long run, over repeated experiments of a similar nature. The remote viewing experiments are no exception. Remember that such events should not replicate with any degree of precision in the short run because of statistical variability, just as we would not expect to always get five heads and five tails if we flip a coin ten times, or see the same batting averages in every game.

The analysis of SRI experiments conducted in 1988 singled out the laboratory remote viewing sessions performed by six "expert" remote viewers, numbers 002, 009, 131, 372, 414 and 504. These six individuals contributed 196 sessions. The resulting effect size was 0.385 (May *et al.*, 1988, p. 13). The SRI analysis does not include information individually by viewer, nor does it include information about how many of the 196 sessions used static versus dynamic targets. One report provided to the review team (May, Lantz and Piantineda, 1994) included an additional experiment conducted after the 1988 review was performed, in which Viewer 009 participated with 40 sessions. The effect size for Viewer 009 for those sessions was 0.363. None of the other five SRI experts were participants.

The same subject identifying numbers were used at SAIC, so we can compare the performance for these individuals at SRI and SAIC. Of the six, three were specifically mentioned as participating in the SAIC remote viewing experiments. As can be seen in Table 2, viewers 009, 131 and 372 all participated

in Experiment 1 and viewers 009 and 372 participated in Experiments 4, 9 and 10 as well.

The overall effect sizes for two of the three, viewers 009 and 372, were very close to the SRI effect size of 0.385 for these subjects, at 0.35 and 0.34, respectively, and the 0.35 effect size for Viewer 009 was very similar to his 0.363 effect size in the report by May, Lantz and Piantineda (1994). Therefore, we see a repeated and, more importantly, hopefully a repeatable level of functioning above chance for these individuals. An effect of this size should be reliable enough to be sustained in any properly conducted experiment with enough trials to obtain the long run statistical replicability required to rule out chance.

It is also important to notice that viewers 009 and 372 did well on the same experiments and poorly on the same experiments. In fact the correlation between their effect sizes across experiments is 0.901, which is very close to a perfect correlation of 1.0. This kind of consistency warrants investigation to determine whether it is the nature of the experiments, a statistical fluke or some methodological problems that led these two individuals to perform so closely to one another. If methodological problems are responsible, then they must be subtle indeed because the methodology was similar for many of the experiments, yet the results were not. For instance, procedures for the sessions with static and dynamic targets in Experiment 1 were almost identical to each other, yet the dynamic targets did not produce evidence of psychic functioning (α -value = 0.50) and the static targets did (p -value = 0.0073). Therefore, a methodological problem would have had to differentially affect results for the two types of targets, even though the assignment of target type was random across sessions.

4.5 Methodological Issues in the Remote Viewing Experiments at SAIC

As noted in Section 2.3, there are a number of methodological considerations needed to perform a careful remote viewing experiment. Information necessary to determine how well each of these were addressed is generally available in the reports, but in some instances I consulted Dr. May for additional information. As an example of how the methodological issues in Section 2.3 were addressed, an explanation will be provided for Experiment 1.

In this experiment the viewers all worked from their homes (in New York, Kansas, California, and Virginia). Dr. Nevin Lantz, who resided in Pennsylvania, was the principal investigator. After each session, viewers faxed their response to Dr. Lantz and mailed the original to SAIC. Upon receipt of the fax, Dr. Lantz mailed the correct answer to the viewer. The viewers were supposed to mail their original responses to SAIC immediately, after faxing them to Dr. Lantz. According to Dr. May, the faxed versions were later compared with the originals to make sure the originals were sent without any changes. Here are how the other methodological issues in Section 2.3 were handled:

- *No one who has knowledge of the specific target should have any contact with the viewer until after the response has been safely secured.* No one involved with the experiment had any contact with the viewers, since they were not in the vicinity of either SAIC or Dr. Lantz's home in Pennsylvania.
- *No one who has knowledge of the specific target or even of whether or not the session was successful should have any contact with the judge until after that task has been completed.* Dr. Lantz and the individual viewers were the only ones who knew the correct answers, but according to Dr. May, they did not have any contact with the judge during the period of this experiment.
- *No one who has knowledge of the specific target should have access to the response until after the judging has been completed.* Again, since only the viewers and Dr. Lantz knew the correct target, and since the responses were mailed to SAIC by the viewers before they received the answers, this condition appears to have been met.
- *Targets and decoys used in judging should be selected using a well-tested randomization device.* This has been standard practice at both SRI and SAIC.
- *Duplicate sets of targets photographs should be used, one during the experiment and one during the judging, so that no cues (like fingerprints) can be inserted onto the target that would help the judge recognize it.* This was done; Dr. Lantz maintained the set used during the experiment while the set used for judging was kept at SAIC in California.
- *The criterion for stopping an experiment should be defined in advance so that it is not called to a halt when the results just happen to be favorable. Generally, that means specifying the number of trials in advance, but some statistical procedures require other stopping rules. The important point is that the rule be defined in advance in such a way that there is no ambiguity about when to stop.* In advance it was decided that each viewer would contribute 40 trials, ten under each of four conditions (all combinations of sender/no sender and static/dynamic). All sessions were completed.
- *Reasons, if any, for excluding data must be defined in advance and followed consistently, and should not be dependent on the data. For example, a rule specifying that a trial could be aborted if the viewer felt ill would be legitimate, but only if the trial was aborted before anyone involved in that decision knew the correct target.* No such reasons were given, nor was there any mention of any sessions being aborted or discarded.
- *Statistical analyses to be used must be planned in advance of collecting the data so that a method most favorable to the data isn't selected post hoc. If multiple methods of analysis are used the corresponding conclusions must recognize that fact.* The standard rank-order judging had

been planned, with results reported separately for each of the four conditions in the experiment for each viewer. Thus, 20 effect sizes were reported, four for each of the five viewers.

4.6 Was Anything Learned at SAIC?

4.6.1 Target Selection. In addition to the question of whether or not psychic functioning is possible, the experiments at SAIC were designed to explore a number of hypotheses. Experiments 1 and 10 were both designed to see if there is a relationship between the "change in visual entropy" in the targets and the remote viewing performance.

Each of the five senses with which we are familiar is a change detector. Our vision is most readily drawn to something that is moving, and in fact if our eyes are kept completely still, we cease to see at all. Similarly, we hear because of moving air, and our attention is drawn to sudden changes in sound levels. Other senses behave similarly. Thus, it is reasonable that if there really is a "psychic sense" then it would follow that same pattern.

Experiments 1 and 10 were designed to test whether or not remote viewing performance would be related to a particular type of change in the target material, namely the "change in visual entropy." A target with a high degree of change would be one in which the colors changed considerably throughout the target. A detailed explanation can be found in the SAIC reports of this experiment, or in the article "Shannon Entropy: A Possible Intrinsic Target Property" by May, Spottiswoode and James, in the *Journal of Parapsychology*, December 1994. It was indeed found that there was a correlation between the change in entropy in the target and the remote viewing quality. This result was initially shown in Experiment 1 and replicated in Experiment 10. A simulation study matching randomly chosen targets to responses showed that this was unlikely to be an artifact of target complexity or other features.

It is worth speculating on what this might mean for determining how psychic functioning works. Physicists are currently grappling with the concept of time, and cannot rule out precognition as being consistent with current understanding. Perhaps it is the case that we do have a psychic sense, much like our other senses, and that it works by scanning the future for possibilities of major change much as our eyes scan the environment for visual change and our ears are responsive to auditory change. That idea is consistent with anecdotal reports of precognition, which are generally concerned with events involving major life change. Laboratory remote viewing may in part work by someone directing the viewer to focus on a particular point in the future, that in which he or she receives the feedback from the experiment. It may also be the case that this same sense can scan the environment in actual time and detect change as well.

Another hypothesis put forth at SAIC was that laboratory remote viewing experiments are most likely to be successful if the pool of potential targets is

neither too narrow nor too wide in terms of the number of possible elements in the target. They called this feature the "target-pool bandwidth" and described it as the number of "differentiable cognitive elements." They reasoned that if the possible target set was too small, the viewer would see the entire set and be unable to distinguish that information from the psychic information. If the set was too broad, the viewer would not have any means for editing an extensive imagination.

Combining these two results would indicate that a good target set would contain targets with high change in visual entropy, but that the set would contain a moderately-sized set of possibilities. The set of 100 *National Geographic* photographs used in the later days at SRI and at SAIC may have inadvertently displayed just those properties.

4.6.2. Remote Staring. Experiment 7, described in Appendix 2, provided results very different from the standard remote viewing work. That experiment was designed to test claims made in the Former Soviet Union and by some researchers in the United States, that individuals could influence the physiology of another individual from a remote location. The study was actually two separate replications of the same experiment, and both replications were successful from a traditional statistical perspective. In other words, it appeared that the physiology of one individual was activated when he or she was being watched by someone in a distant room. If these results are indeed sound, then they may substantiate the folklore indicating that people know when they are being observed from behind.

4.6.3 Enhanced Binary Computer Guessing. Experiment 2 was also very different from the standard remote viewing experiments, although it was still designed to test anomalous cognition. Three subjects attempted to use a statistical enhancement technique to increase the ability to guess forced choice targets with two choices. This clever computer experiment showed that for one subject, guessing was indeed enhanced from a raw rate of just above chance (51.6% instead of 50%) to an enhanced rate of 76 percent. The method was extremely inefficient, and it is difficult to imagine practical uses for this ability, if indeed it exists.

5. External Validation: Replications of Other Experiments

5.1. Conceptual Similarity: Ganzfeld Experiments

While remote viewing has been the primary activity at SRI and SAIC, other researchers have used a similar technique to test for anomalous cognition, called the ganzfeld. As noted in the SAIC Final Report of 29 Sept. 1994, the ganzfeld experiments differ from remote viewing in three fundamental ways. First, a "mild altered state is used," second, senders are [usually] used, so that telepathy is the primary mode, and third, the *receivers* (viewers) do their own judging just after the session, rather than having an independent judge.

The ganzfeld experiments conducted at Psychophysical Research Laborato-

ries (PRL) were already mentioned in Section 3.4. Since the time those results were reported, other laboratories have also been conducting ganzfeld experiments. At the 1995 Annual Meeting of the Parapsychological Association, three replications were reported, all published in the peer-reviewed Proceedings of the conference.

The ganzfeld experiments differ in the preferred method of analysis as well. Rather than using the sum of the ranks across sessions, a simple count is made of how many first places matches resulted from a series. Four rather than five choices are given, so by chance there should be about 25% of the sessions resulting in first place matches.

5.2 *Ganzfeld* Results from Four Laboratories

In publishing the ganzfeld results from PRL, Bem and Honorton (1994) excluded one of the studies from the general analysis for methodological reasons, and found that the remaining studies showed 106 hits out of 329 sessions, for a hit rate of 32.2 percent when 25 percent was expected by chance. The corresponding *p*-value was 0.002. As mentioned earlier, the hallmark of science is replication. This result has now been replicated by three additional laboratories.

Bierman (1995) reported four series of experiments conducted at the University of Amsterdam. Overall, there were 124 sessions and 46 hits, for a hit rate of 37 percent. The hit rates for the four individual experiments were 34.3 percent, 37.5 percent, 40 percent and 36.1 percent, so the results are consistent across his four experiments.

Morris, Dalton, Delanoy and Watt (1995) reported results of 97 sessions conducted at the University of Edinburgh in which there were 32 successes, for a hit rate of 33 percent. They conducted approximately equal numbers of sessions under each of three conditions. In one condition there was a known sender, and in the other two conditions it was randomly determined at the last minute (and unknown to the receiver) that there would either be a sender or not. Hit rates were 34 percent when there was a known sender and when there was no sender, and 28 percent when there was a sender but the receiver did not know whether or not there would be. They did discover post *hoc* that one experimenter was more successful than the other two at achieving successful sessions, but the result was not beyond what would be expected by chance as a post *hoc* observation.

Broughton and Alexander (1995) reported results from 100 sessions at the Institute for Parapsychology in North Carolina. They too found a similar hit rate, with 33 hits out of 100 sessions, or 33 percent hits.

Results from the original ganzfeld work and these three replications are summarized in Table 3, along with the SRI and SAIC remote viewing results. The effect sizes for the ganzfeld replications are based on Cohen's *h*, which is similar in type to the effect size used for the remote viewing data. Both effect

sizes measure the number of standard deviations the results fall above chance, using the standard deviation for a single session.

TABLE 3
Remote Viewing and Ganzfeld Replications

Laboratory	Sessions	Hit Rate	Effect Size
All Remote Viewing at SRI	770	N/A	0.209
All Remote Viewing at SAIC	445	N/A	0.230
PRL, Princeton, NJ	329	32 percent	0.167
University of Amsterdam, Netherlands	124	37 percent	0.261
University of Edinburgh, Scotland	97	33 percent	0.177
Institute for Parapsychology, NC	100	33 percent	0.177

5.3 Conclusions About External Replication

The results shown in Table 3 show that remote viewing has been conceptually replicated across a number of laboratories, by various experimenters and in different cultures. This is a robust effect that, were it not in such an unusual domain, would no longer be questioned by science as a real phenomenon. It is unlikely that methodological problems could account for the remarkable consistency of results shown in Table 3.

6. Is Remote Viewing Useful?

Even if we were all to agree that anomalous cognition is possible, there remains the question of whether or not it would have any practical use for government purposes. The answer to that question is beyond the scope of this report, but some speculations can be made about how to increase the usefulness.

First, it appears that anomalous cognition is to some extent possible in the general population. None of the ganzfeld experiments used exclusively selected subjects. However, it also appears that certain individuals possess more talent than others, and that it is easier to find those individuals than to train people. It also appears to be the case that certain individuals are better at some tasks than others. For instance, Viewer 372 at SAIC appears to have a facility with describing technical sites.

Second, if remote viewing is to be useful, the end users must be trained in what it can do and what it cannot. Given our current level of understanding, it is rarely 100 percent accurate, and there is no reliable way to learn what is accurate and what is not. The same is probably true of most sources of intelligence data.

Third, what is useful for one purpose may not be useful for another. For instance, suppose a remote viewer could describe the setting in which a hostage is being held. That information may not be any use at all to those unfamiliar with the territory, but could be useful to those familiar with it.

7. Conclusions and Recommendations

It is clear to this author that anomalous cognition is possible and has been demonstrated. This conclusion is not based on belief, but rather on commonly accepted scientific criteria. The phenomenon has been replicated in a number of forms across laboratories and cultures. The various experiments in which it has been observed have been different enough that if some subtle methodological problems can explain the results, then there would have to be a different explanation for each type of experiment, yet the impact would have to be similar across experiments and laboratories. If fraud were responsible, similarly, it would require an equivalent amount of fraud on the part of a large number of experimenters or an even larger number of subjects.

What is not so clear is that we have progressed very far in understanding the mechanism for anomalous cognition. Senders do not appear to be necessary at all; feedback of the correct answer may or may not be necessary. Distance in time and space do not seem to be an impediment. Beyond those conclusions, we know very little.

I believe that it would be wasteful of valuable resources to continue to look for proof. No one who has examined all of the data across laboratories, taken as a collective whole, has been able to suggest methodological or statistical problems to explain the ever-increasing and consistent results to date. Resources should be directed to the pertinent questions about how this ability works. I am confident that the questions are no more elusive than any other questions in science dealing with small to medium sized effects, and that if appropriate resources are targeted to appropriate questions, we can have answers within the next decade.

References

- Bem, Daryl J. and Charles Honorton (1994). Does psi exist? Replicable evidence for an anomalous process of information transfer. *Psychological Bulletin*, 115, 4.
- Bierman, Dick J. (1995). *The Amsterdam Ganifeld Series III & IV: Target clip emotionality, effect sizes and openness*. In Proceedings of the 38th Annual Parapsychological Association Convention, 27.
- Broughton, Richard and Cheryl Alexander (1995). *Autoganzfeld II: The first 100 sessions*. In Proceedings of the 38th Annual Parapsychological Association Convention, 53.
- May, Edwin C. (1995). *AC Technical trials: Inspiration for the target entropy concept*. SAIC Technical Report, May 26.
- May, Edwin C., Nevin D. Lantz and Tom Piantineda (1994). *Feedback considerations in anomalous cognition experiments*. Technical Report, November 29.
- May, Edwin C., J. M. Utts, V. V. Trask, W. W. Luke, T. J. Frivold and B. S. Humphrey (1988). *Review of the psychoenergetic research conducted at SRI International (1973-1988)*. SRI International Technical Report, March 1989.

Morris, Robert L., Kathy Dalton, Deborah Delaney and Caroline Watt (1995). *Comparison of the sender/no sender condition in the Ganzfeld*. Proceedings of the 38th Annual Parapsychological Association Convention, 244.

Puthoff, Harold E. and Russell Targ (1975). *Perceptual augmentation techniques: Part two — research report*. Stanford Research Institute Final Report, Dec. 1.

Appendix 1: Effect Size Measure Used With Rank Order Judging

In general, effect sizes measure the number of standard deviation the true population value of interest falls from the value that would be true if chance alone were at work. The standard deviation used is for one subject, trial, etc., rather than being the standard error of the sample statistic used in the hypothesis test.

In rank-order judging, let R be the rank for one trial. If the number of possible choices is N , then we find:

$$E(R) = (N + 1)/2$$

and

$$\text{Var}(R) = (N^2 - 1)/12.$$

Therefore, when $N = 5$, we find $E(R) = 3$ and $\text{Var}(R) = 2$. The effect size is therefore:

$$\text{Effect Size} = (3.0 - \text{Average Rank})/2^{1/2}.$$

Appendix 2: A Brief Description of the SAIC Experiments

Experiments Involving Remote Viewing

There were six experiments involving remote viewing, done for a variety of purposes.

Experiment 1: Target and Sender Dependencies

Purpose: This experiment was designed to test whether or not a sender is necessary for successful remote viewing and whether or not dynamic targets, consisting of short video clips, would result in more successful remote viewing than the standard *National Geographic* photographs used in most of the SRI experiments.

Method: Five experienced remote viewers participated, three of whom (#s 009, 131 and 372) were included in the experienced group at SRI; their identification numbers were carried over to the SAIC experiments. Each viewer worked from his or her home and faxed the results of the sessions to the principal investigator, Nevin Lantz, located in Pennsylvania. Whether the target was static or dynamic and whether or not there was a sender was randomly determined and unknown to the viewer. Upon receiving the fax of the response, Dr. Lantz mailed the correct answer to the viewer. The original response was sent to SAIC in California, where the results were judged by an analyst blind to the correct target. Standard rank-order judging was used.

Since it is not explicitly stated, I asked Dr. May what measures were taken to

make sure the viewer actually mailed the original response to SAIC before receiving the correct answer in the mail. He said that the original faxed responses were compared with the responses received by SAIC to make sure they were the same, and they all were.

Results: Each viewer contributed ten trials under each of the four possible conditions (sender/no sender and static/dynamic target), for a total of 40 trials per viewer. There was a moderate difference (effect size = 0.121, $p = 0.08$) between the static and dynamic targets, with the traditional National Geographic photographs faring better than the dynamic video clips. There was no noticeable difference based on whether or not a sender was involved, supporting the same conclusion reached in the overall analysis of the SRI work. Combined over all conditions and all viewers, the effect size was 0.124 ($p = 0.04$); for the static targets alone it was 0.248 (exact $p = 0.0073$) while for the dynamic targets it was 0.00 ($p = 0.50$).

Discussion: The SAIC staff speculated that the dynamic targets were not successful because the possibilities were too broad. They chose a new set of dynamic targets to be more similar to the static targets and performed another experiment the following year to compare the static targets with the more similar set of dynamic ones. That experiment is described below (Experiment 10.)

Experiment 4: *Enhancing* Detection of AC with Binary Coding

Purpose: This experiment was designed to see if remote viewing could be used to develop a message-sending capability by focusing on the presence or absence of five specific features of a target. The target set was constructed in packets of four, with possible combinations of the absence (0) or presence (1) of each of the five features chosen to correspond to the numbers 00000, 01110, 10101, and 11011. This is standard practice in information theory when trying to send a two digit number (00, 01, 10 or 11); the remaining three bits are used for "error correction." Different sets of five features were used for each of ten target packs.

Method: Five viewers each contributed eight trials, but the same eight targets were used for all five viewers. There was no sender used, and viewers were told that each target would be in a fixed location for one week. They were to spend 15 minutes trying to draw the target, then fax their responses to SAIC in California. The results were blind-judged and the binary features were coded by both the viewers and an independent analyst.

Results: The results were unsuccessful in showing any evidence of psychic functioning. Neither standard rank-order judging nor analysis based on the binary guesses showed any promise that this method works to send messages.

Experiment 5: AC in Lucid Dreams (Baseline)

Purpose: Despite its name, this experiment did not involve lucid dreaming.

Instead, it was used to test three novice remote viewers who were to participate in an experiment involving remote viewing while dreaming. This baseline experiment was designed to see if these individuals would be successful at standard laboratory remote viewing.

Method: For this baseline experiment, each of the three viewers contributed eight trials using a standard protocol common in the SRI era. For each trial, a target was randomly chosen from the set of 100 *National Geographic* targets used at SRI and SAIC. The target was placed on a table (so no sender was used) while the viewer, in another room, was asked to provide a description. The response was later blind-judged by comparing it to the target and four decoys, and providing a rank-ordering of the five choices.

Results: Of the three novice viewers, one obtained a promising effect size of 0.265, although the result was not statistically significant due to the small number of trials (8). Individual results were not provided for the other two viewers, but the overall effect size was reported as 0.088 for the three viewers.

Experiment 6: AC in Lucid Dreams (Pilot)

Purpose: A lucid dream is a dream in which one becomes aware that he or she is dreaming, and can control subsequent events in the dream. This ability has apparently been successfully trained by Dr. Stephen LaBerge of the Lucidity Institute. He was the Principal Investigator for this experiment. The experiment was designed to see if remote viewing could be successfully employed while the viewer was having a lucid dream.

Method: Seven remote viewers were used; four were experienced SAIC remote viewers and three were experienced lucid dreamers from the Lucidity Institute. The latter three were the novice viewers used in Experiment 5. The experienced SAIC remote viewers were given training in lucid dreaming. The number of trials contributed by each viewer could not be fixed in advance because of the difficulty of attaining the lucid dream state. A total of 21 trials were conducted, with the seven viewers contributing anywhere from one to seven trials each. The report did not mention whether or not the stopping criterion was fixed in advance, but according to Dr. May the experiment was designed to proceed for a fixed time period and to include all sessions attained during that time period.

Unlike with standard well-controlled protocols, the viewers were allowed to take the target material home with them. The targets, selected from the standard *National Geographic* pool, were sealed in opaque envelopes with covert threads to detect possible tampering (there were no indications of such tampering). Viewers were instructed to place the targets at bedside and to attempt a lucid dream in which the envelope was opened and the target viewed. Drawings and descriptions were then to be produced upon awakening.

Results: The results were blind-judged using the standard sum of ranks. Since the majority of viewers contributed only one or two trials, analysis by individual viewer would be meaningless. For the 21 trials combined, the effect

size was 0.368 ($p = 0.046$). Information was not provided to differentiate the novice remote viewers from the experienced ones.

Experiment 9: ERD AC Behavior

Purpose: The remote viewing in this experiment was conducted in conjunction with measurement of brain waves using an EEG. The purpose of the experiment was to see whether or not EEG activity would change when the target the person was attempting to describe was briefly displayed on a computer monitor in a distant room. Details of the EEG portion will be explained as experiment 8. Here, we summarize the remote viewing part of the study.

Method: Three experienced remote viewers (#s 009, 372 and 389) participated. Because of the pilot nature of the experiment, the number of trials differed for each viewer based on availability, with viewers 009, 372 and 389 contributing 18, 24 and 28 trials, respectively. Although it is not good protocol to allow an unspecified number of trials, it does not appear that this problem can explain the results of this experiment.

Results: Responses were blind-judged using standard rank-order analysis. The effect sizes for the viewers 009, 372 and 389 were 0.432 ($p = 0.033$), 0.354 ($p = 0.042$) and 0.177 ($p = 0.175$), respectively. The overall effect size was 0.303 ($p = 0.006$).

Experiment 10: Entropy II

Purpose: This experiment was designed as an improved version of Experiment 1. After the unsuccessful showing for the dynamic targets in Experiment 1, the SAIC team speculated that the "target pool bandwidth" defined as the number of "cognitively differentiable elements" in the target pool might be an important factor. If the possible target material was extremely broad, viewers might have trouble filtering out extraneous noise. If the set of possibilities was too small, as in forced choice experiments, the viewer would see all choices at once and would have trouble filtering out that knowledge. An intermediate range of possibilities, too large to be considered all at once, was predicted to be ideal. The standard *National Geographic* pool seemed to fit that range. For this experiment, a pool of dynamic targets was created with a similar "bandwidth." In both Experiments (1 and 10) the researchers predicted that remote viewing success would correlate with the change in visual entropy of the target, as explained in Section 4.6.1.

Method: Four of the five viewers from Experiment 1 were used (#s 009, 372, 389 and 518). They each contributed equal numbers of sessions with static and dynamic targets, with the viewers blind to which trials had which type. Senders were not used, and all sessions were conducted at SAIC in California, unlike Experiment 1 in which the viewers worked at home. Viewer #372 con-

tributed 15 of each type while the others each contributed 10 of each type. Standard rank-order judging was used.

Results: Table 4 shows the results for this experiment. Unlike in Experiment 1, the static and dynamic targets produced identical effect sizes, with both types producing very successful results. The combined effect size for all trials is 0.55, resulting in a z-score of 5.22.

TABLE 4
Results for Experiment 10

Viewer	Static Targets			Dynamic Targets		
	Rank	ES	P	Rank	ES	P
009	2.20	0.565	0.037	1.70	0.919	1.8×10^{-3}
372	1.87	0.801	9.7×10^{-4}	1.93	0.754	1.8×10^{-3}
389	3.10	-0.071	0.589	3.00	0.000	0.500
518	1.90	0.778	7.2×10^{-3}	2.40	0.424	0.091
Total	2.22	0.550	1.1×10^{-5}	2.22	0.550	1.1×10^{-5}

The Other Experiments at SAIC

There were four additional experiments at SAIC, not involving remote viewing. Two of them (experiments 3 and 8) involved trying to measure brain activity related to psychic functioning and will be described briefly. Experiment 3 used a magnetoencephalograph (MEG) to attempt to detect anomalous signals in the brain when a remote stimulus was present. Due to the background noise in the brain measurements and the expected strength of the signal, the experimenters realized too late that they would not be able to detect a signal even if it existed. Experiment 8 utilized an EEG to try to detect the interruption of alpha waves when a remote viewing target was briefly displayed on a computer monitor in another room. The area of the brain tested was that corresponding to visual stimuli. No significant change in alpha was seen.

The remaining two experiments were replications of previous work measuring psychic functioning in areas other than remote viewing. They will be described in detail.

Experiment 2: AC of Binary Targets

Purpose: This experiment attempted to replicate and enhance random number generator experiments conducted at SRI. In these types of experiments a computer randomly selects one of two choices to be the target, denoted as 0 or 1. The internal workings of the computer then rapidly oscillate between 0 and 1 and the subject pushes a mouse button when he or she thinks the internal choice matches the target choice. This process is repeated over many trials. The computer tabulates the results and the experiment is a success if the subject guesses the correct answer more often than would be expected by chance. The purpose is to see if humans can correctly guess computer-selected binary

targets, and hopefully by extension, correctly solve binary choice problems in real situations. If that were to be the case, then real problems could be posed as binary ones (e.g. is the lost child still in this city or not) to narrow down possibilities.

Method: This SAIC experiment was designed to enhance the accuracy of binary guessing by using a statistical technique called sequential analysis. Rather than just one guess for each decision, the subject continues to guess until the computer ascertains that a decision has been reached. The computer keeps track of the number of times zero and one have each been guessed and announces a decision when one of the choices has clearly won out over the other, or when it is clear that it is essentially an ongoing tie. In the latter case, no decision is recorded. Three subjects participated (#s 007, 083 and 531) in this experiment. Subject #531 had been successful in similar experiments at SRI.

Results: Using this method for enhancing the accuracy of the guesses, subject #531, who had been successful in previous similar experiments, was able to achieve 76 correct answers out of 100 tries. This remarkable level of scoring for this type of experiment resulted in an effect size of 0.520 and a z-score of 5.20. The other two subjects did not differ from chance results, with 44 and 49 correct decisions out of 100 or 101. (One subject accidentally contributed an additional trial.)

Although the result for subject 531 is remarkably successful, it does not represent a very efficient method of obtaining the decision. To reach the 100 decisions required a total of 21,337 guesses, i.e. over 200 guesses for each decision. Of the individual guesses, only 51.6 percent were correct, for an effect size of 0.032, similar to other forced choice experiments. Due to the large number of guesses, the corresponding z-score was 4.65. Combined over all three subjects, 56 percent of the 301 decisions were correct and the effect size was 0.123. The combined results were still statistically significant, with $p = 0.017$, as shown in Table 1.

Experiment 7: Remote Observation

Purpose: It is often reported anecdotally that people know when they are being watched. Two experiments were conducted at SAIC to determine whether or not these anecdotes could be supported by a change in physiology when someone is being observed from a distance. The experimental design was essentially the same for the two experiments. This work was a conceptual replication of results reported by researchers in the Former Soviet Union (FSU), the United States and Scotland. The experiments in the FSU were interpreted to mean that the physiology of the recipient was being manipulated by the sender, an effect that if real could have frightening consequences.

Method: The "observee" was seated in a room with a video camera focused on him or her, and with galvanic skin response measurements being recorded. In a distant room the "observer" attempted to influence the physiology of the

observee at randomly spaced time intervals. During those time intervals, an image of the observee appeared on a computer monitor watched by the observer. During "control" periods, the video camera remained focused on the observee but the computer monitor did not display his or her image to the observer. There were 16 "influence" periods randomly interspersed with 16 "control" periods, each of 30 seconds, with blank periods of 0 to 5 seconds inserted to rule out patterns in physiology.

Results: To determine whether or not the galvanic skin response of the observees was activated while they were being watched, the response during the control periods was compared with the response during the "influence" periods for each subject. The results were then averaged across subjects. In both experiments, there was greater activity during the periods of being watched than there was during the control periods. The results were statistically significant in each case ($p = 0.036$ and 0.014) and the effect sizes were similar, at 0.39 and 0.49 . As preplanned, the results were combined, yielding an effect size of 0.39 ($p = 0.005$). As an interesting *post hoc* observation, it was noted that the effect was substantially stronger when the observer and observee were of opposite sexes than when they were of the same sex.

Discussion: This experiment differs from the others conducted at SAIC since it involves interaction between two people rather than one person ascertaining information about the environment or the future. It raises substantially more questions than it answers, since the mechanism for the shift in physiology is unknown. Possibilities range from the idea that the observee was able to know when the computer in the distant room was displaying his or her image, not unlike remote viewing, to the possibility that the observer actually did influence the physiology of the observee. Further experimentation as well as a review of similar past experiments may be able to shed light on this important question.