

Feedback Considerations in Anomalous Cognition Experiments

by

Edwin C. May
Nevin D. Lantz
Science Applications International Corporation
Cognitive Sciences Laboratory
Menlo Park, CA
and
Tom Piantineda
SRI International
Menlo Park, CA

Abstract

In order to determine from what time frame the data from anomalous cognition (AC) originates, we have examined the role of precognition and feedback on the quality of AC. In an otherwise standard AC protocol, we displayed feedback tachistoscopically to receivers. The cognitive awareness of the feedback content was minimal, and 2 of the 8 intensities used for visual display of the feedback were below the visual recognition threshold. We hypothesized a number of possible relationships between feedback intensity and AC quality, including one based on precognition (i.e., the data originated from the future feedback). Four viewers contributed 40 trials each (5 at each of 8 different intensity bands). Using a sum-of-ranks statistic, two viewers produced independently significant evidence of remote viewing (i.e., binomial probability of 2 hits in 4 trials with an event probability of 0.05 is 0.014). None of the data showed a significant positive correlation of feedback intensity with AC quality; one receiver showed a significant negative correlation. This result is discussed with regard to precognition in general and the troublesome unfalsifiability aspect of truly goal-oriented precognition.

Introduction

Beginning in 1986, we conducted a 2-year investigation of the dependency of the quality of anomalous cognition (AC) on the feedback to the participant; hereafter called the receiver.* The experiment was conceptually quite simple, but to address precognitive issues the experiment became technologically complex. Related to the general feedback issue, we were interested in determining from what time frame a receiver gains access to the target.

One model of AC, which is based on precognition, is that the data originate from the future feedback to the receiver (Feinberg, 1975). If we are to understand the process of AC from this perspective, we must examine and understand precognition. One view is that precognition is simply a "clean" methodology for conducting free-response experiments. For the purposes of this paper we have assumed that precognition is real. We cite Honorton and Farrari's (1989) meta-analysis of the forced-choice experiments from 1935 to 1987 as, at least, *prima facie* evidence for the phenomenon. Using 309 studies, they found overall significance ($Z = 11.4, p \leq 2.4 \times 10^{-30}, ES = 0.020 \pm 0.002$), and they examined a variety of variables, including file-drawer and quality considerations.

Targ and Targ (1986) conducted an experiment to explore the relationship between feedback and precognition; do individuals have access to actualized or probable futures? Their evidence suggests that precognition is independent of *a priori* target probabilities. However, in a similar experiment, Radin (1988) subsequently found significant evidence contradicting their result.

In a carefully constructed experiment, Vassy (1986) found that a goal-oriented model of precognition did not fit his data; however, our analysis (May, Utts, Spottiswoode, and James, 1994) demonstrated that the timing parameters that Vassy used in his experiment could not exclude a goal-oriented interpretation of the results.

We will discuss the apparent contradictory results of these experiments in the Discussion Section of this paper. We note that neither the Targ and Targ nor the Vassy studies explored the functional relationship between feedback and AC. Nonetheless, because of their more recent investigations and because of the importance of the question about feedback, we are presenting an updated and more complete version of our 1986 experiment than appeared earlier (May, 1993).

Conceptual Approach

In a schematic protocol for a typical AC trial, a receiver and monitor (i.e., interviewer) are sequestered in a laboratory. An assistant randomly selects a target from a pre-defined set (i.e., target pool), and at a predetermined time, the receiver attempts to describe the target. At the end of the session, the AC data are secured, and the intended target is shown to the receiver as feedback. Normally the feedback includes a presentation of the target and involves a complete debriefing of the session experience.

In our experiment, we attempted to inhibit as many potentially open precognition channels as possible so that we could assess the impact of the visual intensity of the feedback display on the quality of the AC. Thus, we eliminated all debriefing of the AC experience during the feedback periods and presented

* The Cognitive Sciences Laboratory has adopted the term *anomalous mental phenomena* instead of the more widely known *psi*. Likewise, we use the terms *anomalous cognition* and *anomalous perturbation* for ESP and PK, respectively. We have done so because we believe that these terms are more naturally descriptive of the observables and are neutral in that they do not imply mechanisms. These new terms will be used throughout this paper.

visual feedback tachistoscopically. The display intensities for the feedback presentation varied from zero to a level that just exceeded the visual recognition threshold (i.e., some details can be recognized). Even the strongest display intensity was insufficient to provide a "satisfying" study of the target material.

To lessen the impact of other potential precognitive leakage paths, extreme care was taken in order to insure that the receiver was the only individual who was simultaneously aware of both the intended target and the response. That is, during the course of the experiment and to date, no individual except the receiver has been consciously aware of the target that matches any given response.

Under these conditions we hypothesized four potential outcomes, which are displayed schematically in Figure 1. Under the null hypothesis of no AC, the quality is low and does not depend upon the intensity of the feedback. If precognition of the feedback is the underlying mechanism, then we would expect a linear relationship; the more information in the feedback, the higher the quality of the AC. That is, the more information in the receiver's future, the more AC in the session. Suppose, however, that the information is acquired exclusively in "real-time." That is, none of the information arises from the future feedback. Since, by definition, there is no dependency on the feedback, we might expect significant evidence for AC like the one labeled "Real-time." A combination of both mechanisms is also shown. For example, if the feedback is not cognitive (i.e., subliminal), there is little information in the feedback so that the only open channel is real-time; whereas, the opposite might be true for large amounts of feedback information.




Figure 1. Idealized Relationships Between AC Quality and Intensity of Feedback

In this discussion, we have assumed that the feedback content is proportional to the *cognitive* awareness of the visual feedback in our proposed experimental outcomes. Based on this assumption, we used the amount of information available (i.e., the visual intensity) as feedback as the independent variable in an otherwise straight forward AC experiment.

Experiment Details

In this section, we describe the methodological and technical details of an experiment to determine the degree to which the quality of the AC depends upon the intensity of the visual feedback.

Tachistoscope Calibration

The crucial independent variable is the amount of visual feedback displayed to the receiver. The magnitude of the stimulus is directly proportional to the duration of the receiver's exposure for a given level of luminance. In a calibration experiment, volunteers were presented with slides and asked to say when they were aware of the presentation. We manipulated the magnitude of the stimuli from zero to a value where the volunteer could recognize the gestalt of a scene. Each stimulus slide was presented for 50 milliseconds (ms), and the magnitude of the stimulus information was adjusted by attenuating the luminance of the slides over a range of two logarithmic units. In adjusting the magnitude, we relied upon Bloch's Law, which says that for presentation time shorter than about 100 ms, the product of time and intensity is constant (Marks, 1975). Thus, varying the luminance of the stimuli is equivalent to varying its duration. In addition each slide was appropriately masked with random intensity patterns.

For the luminance calibration, a two-field projection tachistoscope (i.e., Gerbrands G1170) was loaded with 80 color photographic 35 mm slides (5 opaque and 75 having various luminance contrasts) of natural and man-made scenes, which were randomly chosen from our larger pool of 200 photographs from *Natural Geographic* magazines (May, Utts, Humphrey, Luke, Frivold, and Trask, 1990). We varied the luminance contrast of the slides by duplicating them at one of twelve f-stops (including zero) to provide a target pool having variations in the intensity covering two logarithmic units. The contrast in luminance for each slide, which may be considered to be the ratio of the brightest to the darkest part of the slide, was further attenuated in the calibration trials so that some of the slides were above and others below the observer's detection threshold. To avoid any possible audio cue to the presentations, the individuals were presented with white noise in a large, comfortable head phone set, and the tachistoscope was located in an adjacent room.

The calibration slides were tachistoscopically back-projected onto a 14-inch-square frosted glass window. Figure 2 shows the floor plan that was used both for the calibration and the AC measurements. During the calibration experiments, the window cover was fully open; however, during the AC sessions, the cover was positioned to prevent the monitor from seeing the frosted glass screen.

The tachistoscope was programmed to present each calibration stimulus in numerical order for 50 ms, followed by a 5-second pause during which the next slide was cycled into position. Slides were attenuated by projecting them through a pair of plane polarizers: one fixed and the other variable. The luminance of the projected image varied as the cosine of the angle between the two polarizers.

Visual detection, however, is not related to luminance alone. For a fixed luminance, scenes with different contrasts will be detected with differing probabilities. At the same luminance, for example, a photograph of a checkerboard will be easier to detect than a photograph of a mountain cabin in a snow storm. Thus each of the candidate target slides had to be calibrated with human observers in order to determine the empirical relationship between detectability and luminance.

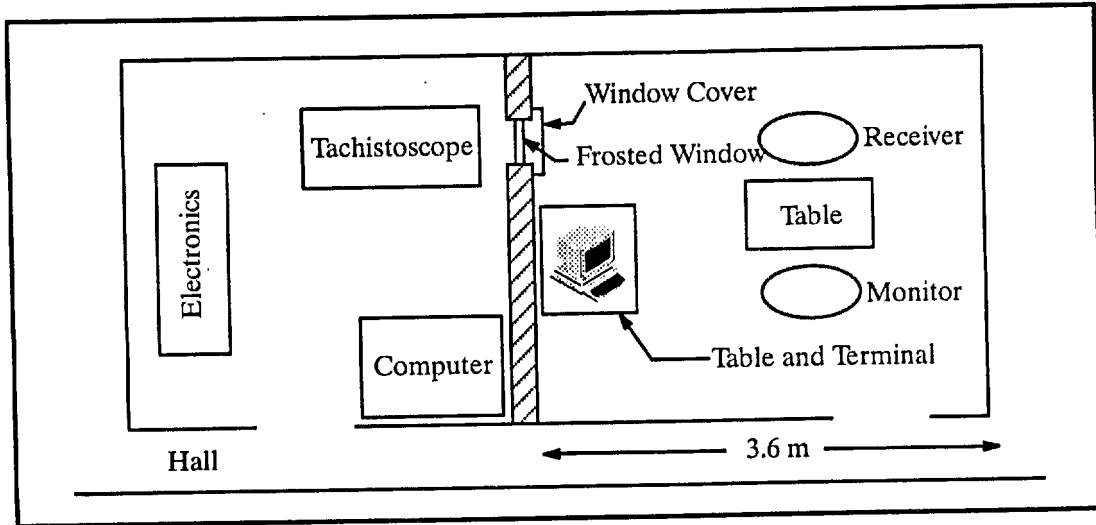


Figure 2. Schematic Floor Plan.

Two naive female volunteers participated in the calibration. A complete data set was obtained from one volunteer, and data trends were confirmed by the second volunteer.

The calibration procedures were as follows. The volunteer was seated approximately three feet from the projection screen, which was positioned at eye level in the wall between the room in which the apparatus was housed and the room in which the volunteer sat (see Figure 2). The volunteer was permitted to view the screen and the other contents of the room freely for several (i.e., more than five) minutes to ensure that she adapted to the ambient illumination level. To screen the sound of the tachistoscope, the volunteer listened to white noise through earphones. The response was registered by a foot switch that the volunteer pressed to indicate detection of the stimulus slide. In a typical session, the variable polarizer was set at a predetermined value and each of the 80 slides was presented 5 times. Two sessions were conducted at each polarizer setting, providing ten data points per slide per polarizer setting. An alternative procedure was used when the variable polarizer was set near one of the extremes of the experimental range.* To reduce the tedium, only those slides near the detection threshold were presented.

Each time a new slide was presented, the volunteer reported whether the presentation was detected (i.e., sensed at all, regardless of target content). Counters recorded whether a particular slide was detected as well as the proportion of slides detected. From these records, a psychometric function was generated relating the proportion of the time each slide was detected to the contrast in luminance for that slide. This function, which relates the contrast in luminance for the slide to its detection threshold, is an index of the the detectability of the geographic scene depicted in the slide. By using this psychometric function, it was possible to specify not only which slides are subliminal (i.e., never detected), but also how far above or below the detection threshold each slide lies.

Figure 3 shows a sample of the psychometric curves generated from these data. Six of the 80 slides are shown by plotting the probability of detecting a given slide as a function of the variable polarizer setting. We generated a set of curves similar to the ones shown in Figure 3 for all target slides.

* With the extreme settings, the volunteer saw nearly all of the slides or very few of them.




Figure 3. Degrees of Polarization (Scaled for Equal Luminance Intervals)

Normally, data would be collected from a large sample of individuals in order to arrive at an average function, but in this experiment, data from two persons were sufficient for several reasons. First, pilot studies indicated that interperson variability of stimulus slide detection was quite low. Second, to collapse interperson variability even further, we generated a steep psychometric curve by sampling the abscissae coarsely. For example, if we sampled stimulus slide contrast at only two values—zero and 100 percent contrasts—all observers would respond identically, thus eliminating interperson variation if the observers possessed normal or corrected vision. In this study, we sampled stimulus target contrasts at intervals that were found in pilot studies to produce low interperson variability. Finally, for the purposes of this study, interperson variability was not significant because it only shifts the psychometric function along the abscissa by some unknown amount without changing the shape of the function. Thus, interperson variability could only result in an erroneous estimate of feedback magnitude. While these errors may influence the intercept of the function relating the dependent variable (i.e., AC performance) to feedback magnitude, the slope of the performance versus magnitude of the feedback is independent of these errors.

Figure 4 shows an idealized relationship between visual detectability and luminance for this experiment. To lessen the sensitivity to individual differences in perceiving the feedback, we required that the psychometric curve rise rapidly through the 50% detection level. We identified eight levels of feedback intensity that would be used as the values of the independent variable in the experiment. Two were subliminal (i.e., 0.0 detection probability), and one of these was completely opaque; one was nominally at threshold; three were at 25, 50 and 75 percent detection threshold, respectively; and two were above recognition threshold (i.e., 100 percent). These points are shown as arrows in Figure 4.




Figure 4. Ideal Psychometric Function

The following example illustrates the procedure that was used for all target slides to select those that met the criteria shown in Figure 4. We notice in Figure 3, that one of the slides is 50% detectable when the polarizer was set at 50° , which is one of the detection bands required as illustrated in Figure 4.

Suppose we wished to include the candidate slide, which was detected 50% of the time when the polarizer was set at 87.5° , in the 50% detection group. We notice a difference of 37.5° in the polarizer setting between this candidate and the original slide. To account for this difference, we compute $\cos^2(37.5^{\circ}) = 0.629$. When we then reduce the exposure of the candidate slide by this factor, it will also be detected only 50 percent of the time when the polarizer is set at 50 degrees.

In actual practice, we could only change exposure by an integral number of f-stops, so changes of intensity were always by factors of 2. Working backward, however, the position of all the target slides on the curve shown in Figure 4 could be determined. Furthermore, all the slides were grouped into bands as close as possible to the arrows shown in Figure 4.

Experiment Protocol

Forty targets, which were randomly selected from the original 80 used in the calibration, were prepared into eight intensity groups of five targets each using the calibration data and technique described above. Each intensity group represented the cognitive awareness that each receiver would experience (on the average) from the feedback. The top two intensities were sufficient to recognize major features, but insufficient to discern details. By definition, those below the visual recognition threshold could not be cognitively sensed, and one group constituted no feedback, whatsoever (i.e., zero luminance).

To attempt to maintain some control over precognitively available "answers," we arranged that at no future time would a response be knowingly compared to its intended target. For a complete answer to exist in the future, three pieces of information are needed: (1) the target, (2) the response, and (3) the comparison between them. The target system was prepared by individuals who had no access to the responses, and the AC monitor, the assistant, and the receivers had no access to the targets. Finally, the analysts were never informed which were the correct results on a trial-by-trial basis. Thus in the experiment and forever after only the receiver has had access to all three pieces of information.

The slide tray in the tachistoscope was controlled by a computer (Sun Microsystem 3-160) in such a way that all participants were blind to target selection during a trial. To avoid cuing, for example, the tray always began and ended in the zero position. When the computer moved the tray, an independent electrical unit, which could be accessed by the computer, counted the tray steps to assure that the intended target was displayed at the correct time. After the session, the computer repositioned the tray to zero.

Two very experienced receivers (receivers 009 and 177), one moderately experienced receiver (105), and one novice (137) each contributed 40 trials (five at each of the eight intensity levels) at an average rate of about five trials per week. All receivers except 009 possessed normal vision; receiver 009 is color blind. The receivers were all in their thirties, and 009 and 177 were male and 105 and 137 were female. All receivers were completely informed about the nature of the experiment and were asked to sign a consent form indicating that they were also aware of its potential risks. The protocol had been approved by an appropriate Institutional Review Board.

All receivers believed strongly that feedback was not necessary for success, but they were uncertain about the degree to which feedback might contribute to success.

A random order of intensities of the feedback was determined once by a pseudorandom number generator, which was based on a shift-register algorithm by Kendell and has been shown to meet the general criteria for "randomness" (Lewis, 1975), and differently for each receiver prior to the start of the receiver's first trial. Once the order had been set, the trials cycled through the list of intensities until the 40 trials were complete. The sequence of events for each trial was as follows:

- (1) A monitor and a receiver entered a laboratory that contained a table, two chairs, a computer terminal, and a covered 14-inch-square frosted glass window. The window served as a projection screen for the tachistoscope in the adjacent laboratory. See Figure 2.
- (2) The monitor initiated an automatic target selection program on the terminal.
- (3) Using the standard feedback shift-register algorithm, which was seeded from the system clock, the computer randomly selected (with replacement) a target from within the set of five for the given intensity, stepped the slide tray to that target, and notified the monitor that the trial could begin. Because of the closed tachistoscope shutters, no illumination of the slide was present on the frosted screen, and the time to position the slide tray was always set to the time that was required to reach the most distant slide, regardless of the target slide's actual position. Thus, there was no time cuing.
- (4) The monitor indicated that the trial should begin. For the next 10-15 minutes, the receiver drew and/or wrote responses to the intended target.
- (5) At the conclusion of the session, the monitor collected the response, and the receiver opened the screen cover in such a way as to shield the monitor from the feedback material.
- (6) When the receiver was ready, he or she pressed a button that initiated a single tachistoscope display of the target. One and only one, 50-ms display appeared on the translucent window screen. (Electronics prevented the receiver from receiving more feedback after the first button press.) The monitor was instructed *not* to discuss the experience with the receiver in any way at any time.
- (7) The monitor ended the session, and notified the control program from the computer terminal. After the computer had returned the slide tray to zero, then and only then, did the monitor and receiver leave the room. All target data were preserved in a computer file.

Data Analysis

The rank-order analysis was used in this experiment. Using cluster analysis, all 200 targets had previously been assigned to relatively orthogonal clusters of nearly equal number of similar targets. An assistant

prepared packages (one for each receiver) consisting of all the responses randomly ordered. Next, the assistant computer-generated a list (ordered by target number) of seven targets for each response consisting of the actual target and six decoys (a different set of seven for each response). The decoy clusters were chosen randomly, but uniquely, from the complete set minus the one from which that actual target was selected. That is, no two decoys were chosen from the same cluster. Once a cluster was selected, the decoy was randomly selected from within the cluster.

The response material and the target/decoys set of seven photographs (i.e., one target, six decoys) were presented to two analysts for judging. The analysts arrived at a consensus to rank order each set of seven targets for each response in accordance with the best to the worst response/target match. For each receiver, a sum-of-ranks statistic was computed for the sessions. In addition, the data were plotted as AC quality (i.e., eight minus the assigned rank) versus the feedback intensity.

Results

Table 1 shows the average rank (<Rank>), effect size (ES), Z-score, and its associated *p*-value for each receiver in this experiment.*

Table 1.

Results of the Tachistoscope Feedback Experiment

Receiver	<Rank>	ES	Z	P-Value
177	2.600	0.700	4.43	4.43×10^{-6}
009	3.275	0.363	2.29	0.011
137	3.975	0.013	0.08	0.468
105	4.550	-0.275	-1.74	0.959
Total	3.600	0.200	2.53	0.004

The effects size was computed from:

$$ES = \frac{\bar{R}_e - \bar{R}_o}{\sqrt{\frac{N^2 - 1}{12}}}$$

where *N* is the number of possible ranks, which was seven in this experiment; \bar{R}_e and \bar{R}_o were the expected and observed average rank. The Z-score was computed from $ES \times \sqrt{m}$. For the individual receivers, $m = 40$, and for the total $m = 160$.

Receivers 177 and 009 produced independently significant results (1-tailed). We can combine data for all receivers in many ways, but the most conservative is a binomial calculation assuming an event probability of 0.05. Two significant trials in four attempts corresponds to an exact *p*-value of 0.014. The totals shown in Table 1, were computed from the total sum of ranks. The important point, however is that this experiment produced strong evidence for AC.

* The numbers in Table 1, differ slightly from the table on page 137 of May (1993). The results here were checked against the recorded sum of ranks from which the average rank was computed. We believe these results to be correct.

Figure 5 shows AC quality plotted against the intensity of the feedback for the four receivers. The lowest possible quality was one and the maximum was seven (i.e., 8-rank). Shown also is the least squares best fit line.




Figure 5. Combined Data: Quality vs Feedback Luminance.

We computed a Spearman's ρ correlation for each of the receivers and for the combined data. Table 2 shows the results of these calculations.

Table 2.

Spearman's Rank Correlations for Quality with Feedback

Receiver	ρ	t	df	P-Value
177	0.112	0.695	38	0.245
009	-0.297	-1.917	38	0.969
137	0.065	0.402	38	0.345
105	0.150	0.935	38	0.178
Total	0.004	0.050	158	0.480

To place these correlations in perspective, we provide a power analysis for 40 trials. One-tailed correlations of 0.26 and 0.48 would lead to 50% and 80% chance, respectively, of reaching statistical significance (Cohen, 1977).

Figure 6 shows the quality versus feedback luminance for the two independently significant receivers (i.e., 177 and 009). As seen in Table 2, receiver 177 demonstrated a slightly positive correlation with the intensity of the feedback, while receiver 009 showed a significantly negative correlation.

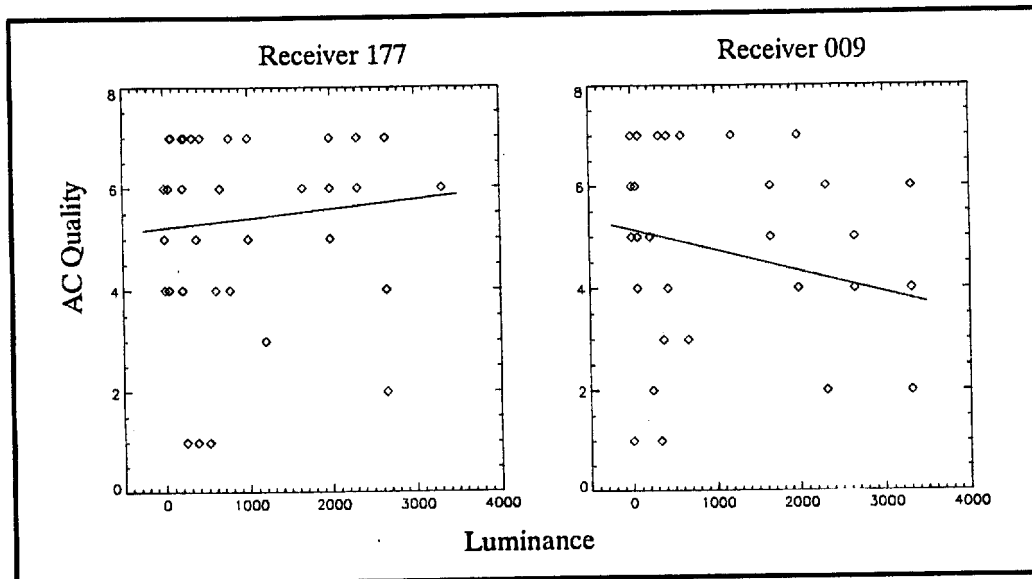


Figure 6. Correlations for the two Significant Receivers.

Discussion

The relationship that is easiest to understand is hypothesis 1 in Figure 1 (i.e., increased performance with increased feedback intensity). We did not observe any such positive correlation overall or with either of the significant receivers. Receiver 009 produced a significant negative correlation between performance and feedback, and at this time, we have no explanation for this result.

This lack of positive correlation, in conjunction with significant evidence of *AC*, complicates the interpretation considerably. The most obvious conclusion is that the receivers obtained their data in real time and not from the feedback. However, because of the conceptual difficulties outlined in the introduction of this paper, we must consider alternative explanations.

One alternative is that the absolute position of the psychometric curve is important. Feedback might be related to the cognitive experience but only at large values of luminance contrast. If this is true, then we might not expect functional dependence of *AC* quality on the feedback in this particular experiment.

An important alternative, however, is that precognition may be exclusively goal oriented, and thus, we are faced with the unfalsifiability issue. We might not ever be able to interpret process-oriented experiments if this aspect of precognition is true. Therefore, the question of from what time frame does *AC* data originate remains unanswered at the present.

In response to a question from Russell Targ in June 1994, we examined, *post hoc*, the data when *no* feedback was presented (i.e., 0.0 intensity). Table 3 shows the effect size for all the trials (*ES_t*), zero-intensity average rank (*<Rank>*), zero-intensity effect size (*ES*), *Z*-score, and its associated *p*-value for each receiver.

Table 3.

Results for Zero Intensity Feedback

Receiver	ES _t	<Rank>	ES	Z	P-Value
177	0.700	2.60	0.700	1.56	0.059
009	0.363	3.00	0.500	1.12	0.132
137	0.013	6.00	-1.000	-2.24	0.987
105	-0.275	5.00	-0.500	-1.12	0.868
Total	0.200	4.15	-0.075	-0.34	0.631

From these results it appears that the most experienced receivers do not require feedback on a trial-by-trial basis; receivers 177 and 009 maintain or increase the effect sizes from their sets as a whole. The less experienced receivers, however, appear to perform considerably below their overall contributions. As Targ and Targ (1986) also found, trial-by-trial feedback is *not* necessary for significant AC.

One of us (May) has been exploring the nature of precognition and suggests that precognition experiments are relatively easy to construct and nearly impossible to understand. For example, if precognition is goal-oriented (i.e., individuals can "peek" into the answer book), then process-oriented experiments are difficult, if not impossible, to interpret. It is always possible to include *all* the complexity of any given experiment into a "black box" and consider the final result as the "answer book." Thus the intervening complexities are simply not apropos.

One obvious problem with this perspective is that goal-oriented precognition is *not* falsifiable—an unacceptable circumstance in science. For any experiment to be valid, there must be a result. No matter how cleverly the result is "hidden" from various participants in the experiment or how contrived the protocol, as long as a result will eventually exist, truly goal-oriented precognition is *always* possible.

The apparent contradiction among the experimental results from Targ and Targ (1986), Radin (1988), and May et. al's (1994) computation of Vassy's results (1986), could be a manifestation of true goal-orientation. It is only the Radin study that suggests that *a priori* target probabilities are important. Radin, however, did not consider *global* goal-orientation in his discussion. In his experiment, a pseudo random number generator determined which face of a computer-die would be *a priori* biased. Since the seed for that decision was taken when "As soon as the cursor is moved at least one pixel (~.3 mm), ..., " (Radin, 1988, page 195) all Radin, the subject, had to do was to "peek" into his future for the trial and compute when to move the cursor to select the seed to correspond to a biased face that did *not* land face up. If this were the case, then we could not interpret Radin's significant precognition results as evidence for a dependency upon *a priori* futures.

There is another technical problem concerning precognition. From a physics point of view, the present does not exist. It is a single geometric point in space-time. (The present is an interval of time with an infinitesimally short duration.) All other times are either in the past or in the future, and thus the definition of real-time AC is problematical.

In an *AC* experiment, what constitutes the “answer book” in the future? The most direct candidate is the reporting of the target material to the receiver after the session (i.e., feedback). It is arguable that providing no feedback to the receiver precludes precognition as an explanation for successful *AC* experiments. Since there are no examples in physics that are fundamentally anthropocentric, and if information from a future time is available to anyone, then most certainly it is available to everyone. The implication is that receivers could obtain information regardless of the feedback parameters—including no feedback at all.

Acknowledgements

We would like to thank Beverly Humphrey for her valuable contribution as an analyst and Thane Frievo for creating the tricky computer control code. Thanks also to the tireless effort of the remainder of the Cognitive Sciences Laboratory for target selection and technical administration.

REFERENCES

- Cohen, J. (1977). *Statistical Power Analysis for the Behavioral Sciences*, Second Edition. Lawrence Erlbaum Associates, Hillsdale, New Jersey. Page 86.
- Feinberg, G. (1975). Precognition—A memory of things future? *Quantum Physics and Parapsychology*, Proceedings of an International Conference held in Geneva, Switzerland, August 26-27, 1974. Laura Oteri, Editor. Parapsychology Foundation, Inc., New York, 54-73.
- Honorton, C. and Ferrari, D.C. (1989). 'Future Telling:' A meta-analysis of forced-choice precognition experiments, 1935-1987. *Journal of Parapsychology*, 53, 281-308.
- Lewis, T. G. (1975). *Distribution Sampling for Computer Simulation*. Lexington, MA: Lexington Books.
- May, E. C. (1993). Technology: A mixed blessing for modern psi research. *Psi Research Methodology: A Re-Examination*, Proceedings of an International Conference held in Chapel Hill, NC, October 29-30, 1988. Lisette Coly and Joanne D. S. McMahon, Editors. Parapsychology Foundation, Inc., New York, 128-148.
- May, E. C., Utts, J. M., Humphrey, B. S., Luke, W. L. W., Frivold, T. J., and Trask, V. V. (1990). Advances in remote-viewing analysis. *Journal of Parapsychology*, 54, 193-228.
- May, E. C., Utts, J. M., Spottiswoode, S. J. and James, C. L. (1994). Applications of decision augmentation theory. Submitted to the *Journal of Parapsychology*.
- Radin, D.I. (1988). Effects of a priori probability on psi perception: Does precognition predict actual or probable futures? *Journal of Parapsychology*, 52, 187-212.
- Targ, E. and Targ, R. (1986). Paranormal perception as a function of target probabilities. *Journal of Parapsychology*, 50, 17-28.
- Vassy, Z. (1986). Complexity dependence in precognition. *Journal of Parapsychology*, 50, 235-270.