13 Radin, D. I., May, E. C., and Thomson, M. J., Psi experiments with random number generators: meta-analysis part 1. SRI International, Menlo Park, CA 1985.
14 Schmidt, H., A PK test with electronic equipment. J. Parapsych. *34* (1970) 175–181.
15 Targ, R., and Puthoff, H. E., Information transmission under conditions of sensory shielding. Nature *252* (1974) 602–607.
16 The Society for Psychical Research: objects of the society. Proc. Soc. psych. Res. *1* (1882–83) 3–6.

0014-4754/88/040315-08$1.50 + 0.20/0
© Birkhäuser Verlag Basel, 1988

Remote viewing

C. Scott

60 Highgate Hill, London N19 5NQ (England)

Summary. Remote viewing is the supposed faculty which enables a percipient, sited in a closed room, to describe the perceptions of a remote agent visiting an unknown target site. To provide convincing demonstration of such a faculty poses a range of experimental and practical problems, especially if feedback to the percipient is allowed after each trial. The precautions needed are elaborate and troublesome; many potential loopholes have to be plugged and there will be strong temptations to relax standards, requiring exceptional discipline and dedication by the experimenters. Most reports of remote viewing experiments are rather superficial and do not permit assessment of the experimental procedures with confidence; in many cases there is clear evidence of particular loopholes left unclosed. Any serious appraisal of the evidence would have to go beyond the reports. Meanwhile the published evidence is far from compelling, and certainly insufficient to justify overthrow of well-established scientific principles.
Key words. Remote viewing; ESP; feedback; data selection; bias; fraud; statistics; methodology.

Introduction

Parapsychology has its fashions and, like all fashions, they tend to move in cycles. The hypnotists of the late 19th century devoted much attention to a phenomenon they called 'travelling clairvoyance': the hypnotic subject was asked to send his mind to a distant place and describe what he saw. It was widely believed that the state of hypnosis conferred extrasensory powers and this was said to be demonstrated when the subject reported veridical information about the location visited which he could not have known by normal means[4].

In the mid-1970s this idea, without the hypnosis, was resuscitated by two California research workers, Targ and Puthoff, under the name of 'remote viewing' (sometimes called 'remote perception'). Targ and Puthoff[18, 19] put their subject, or 'percipient', in their laboratory with a tape recorder while an 'agent' was sent to one or more randomly selected locations. At a pre-arranged time the agent viewed the chosen location and the percipient recorded his impressions. The recordings were transcribed, the experiment was repeated for many different locations, and judges were presented with the set of transcripts together with the list of locations, and asked to match one against the other. Striking successes were claimed. In a short time remote viewing jumped to the forefront of the armoury of methods used by parapsychologists to demonstrate ESP.

New fashions do not arise simply by chance. One of the attractions of remote viewing was its fund-raising potentiality. If a percipient sitting in a laboratory in California can send his mind anywhere in the world, then obviously he can send it also to Moscow and visit the most secret files in the Kremlin. Some researchers were not slow to point this out to the Department of Defense where the thought duly stimulated the issuance of some lucrative research contracts.

Another practical advantage of the method is the ease with which one can demonstrate face validity. A single transcript typically includes scores of statements, while on the other side a target location contains a great number of objects, shapes, sights, sounds and associations. The opportunity to find a striking correspondence somewhere among such a mass of material is large – and more important, much larger than it seems to be to the casual inquirer. Marks and Kammann[9] describe the phenomenon they call 'subjective validation', which leads the observer to seek correspondences between the transcript and the target and when he finds them (as, almost inevitably, he will) to overlook the size of the parent set of comparisons from which the successful correspondences have been selected. Unlike their Victorian forbears, modern parapsychologists are not so naive as to believe that the mere citation of a selection of striking resemblances between statements in the transcript and features of the target provides an adequate demonstration of the reality of remote viewing. They admit that, for this purpose, only a blind matching method is scientifically acceptable. Nevertheless, once this admission has been made, they feel entitled to go ahead with just such a presentation of selected correspondences by way of examples. In doing so, no author has yet considered it necessary to attempt to give an estimate of the size of the parent set of comparisons. The effect on the reader (or the viewer – see particularly the BBC's *Horizon* programme of 1983, 'The Case of ESP') is overwhelming: taken out of context, the hits appear to outstrip any possibility of chance coincidence. Thus the face validity of the remote viewing technique is a valuable feature for the experimenter wishing to gain the support of a non-specialist audience – a category which generally includes TV producers and some Defense Department project evaluators. It also helps to sustain the morale of those working on the project, in sharp contrast to the older card-guessing techniques which are notoriously boring for those involved.

Setting aside the anecdotal presentation of selected hits, how exactly can the reality of remote viewing be demonstrated?

Selection and definition of targets

First, it is clear that the agent must not be free to choose the target location according to his whim. It has been known to psychologists for a long time that seemingly random choices are in reality influenced by a variety of response preferences which may be shared by many or all respondents. If targets were chosen at the whim of the agent, any coincidence observed between the target and the subject's 'perception' could be attributed to common response preferences be-

tween the two persons concerned. It is for this reason that parapsychologists have accepted for many years the necessity of random target selection in their experiments. (Despite this broad concensus, the extensive remote viewing data bank reported by Dunne, Jahn and Nelson [5] includes 45% of trials in which the agent made his own choice of target.) Unfortunately while this principle is widely accepted, its application in the context of a remote viewing experiment poses problems which are often overlooked. The choice of target location may be determined by a rigorous random procedure but the target location often does not determine the target itself. In most experiments the agent has been directed to visit a specific place (for example, the Adler Planetarium in Chicago in one of the experiments reported by Dunne et al. [5]) but, having reached this location, he is free – indeed encouraged – to wander around for the duration of the experiment, typically 15 minutes, taking notes on points of interest, subjective reactions, impressions, feelings and observations and in many cases making sketches. All of this material may be included in the package of target material given to the judge for comparison with the transcript. In some cases photographs have been added later to this material. Thus the target is not the Adler Planetarium but the sum total of the material submitted for judging. The principle of random target selection is clearly being violated here. For example, on a rainy day the agent is likely to stay mainly indoors; on a sunny day he will go outside. The percipient is equally aware of the weather and may adjust his responses in the same way. Thus, shared response preferences may come into play to increase the probability of a hit.

We have mentioned above that a photograph of the target area may be appended to help the judge. If all such photographs are taken before the experiment and appended to the corresponding targets in the target pool there can be no objection. But if, as has happened in some experiments, the photograph is taken after the experiment (worse, by someone who has read the percipient's transcript) then the double-blind character of the judging is seriously compromised. It is all too easy for the photographer to select, perhaps even unwittingly, a view or an angle which draws attention to a feature of the transcript.

Similarly in some experiments the judge is expected to visit the target location and the experimenter therefore prepares instructions to tell him how to get there. These may include the angle of approach and may define more exactly the scope of the location. Again, such material may be biased in such a way as to bring out points of agreement with the percipient's transcript. Hence instructions of this kind must be prepared in advance of the experiment to preserve the double blind feature.

In the case mentioned above, a favourable bias can be introduced by adding materials to the target package. More subtly, bias can also be created by subtraction of such material when it disagrees with the transcript. In some experiments, photographs have been included for some targets but not for others, suggesting the possibility that biases of this kind have been at work.

The only safe precaution is to insist that the complete package of target material be prepared and frozen before the experimental series begins. Neither the agent nor the experimenter should be allowed to add to it once the percipient has started making his description. Each experimental session shoud be preceded by a rigorously random selection of one unit from among this target pool.

Controlling the percipient's response transcripts

Turning now to the other half of the experiment, what precautions are needed in relation to the material produced by the percipient? Obviously the percipient must be free to say anything he wishes. The only requirement is that everything he says, including any drawing he may make, should be included without distortion in the transcript. Again, this apparently simple condition is not always easily attained. The first step is to transcribe the recording. The initial draft is likely to be messy: there will be queries about unclear phrases, incoherencies, interruptions. The script may be excessively long and wandering. The temptation to edit it will be great. In some cases translation may be necessary; here editing of a kind becomes inevitable [16]. Moreover the subject's drawings have to be fitted in. In most cases the need for a second draft will seem evident. But once the principle is accepted that a transcript can be edited we have to raise the same questions about bias as those we raised in the case of the target material. Specifically, has anything in the transcript been added, subtracted or modified, and if so when, by whom and why? In Bisaha and Dunne [2], for example, certain drawings made by percipients are actually mentioned in the transcript but are not included in it. The only safe procedure is to have the person responsible for editing ignorant of the choice of target. In practice this is not easily arranged: the process of initial transcription is likely to take at least a day and the editor is likely to be one of the principal experimenters, the other having been responsible for target selection. Thus these experimenters must not communicate for a day. If sessions are carried out each day this leaves them little or no time to communicate during the whole experiment, a constraint which they are unlikely to respect in practice. A safer and more practical policy would be to ban editing altogether, leaving the preparation of the transcript entirely to the typist (though checked, ultimately, by the experimenter). But a further complication arises if, as is often the practice, the experimental procedure calls for feedback to the percipient after each session. The percipient then knows the earlier targets and may refer to this knowledge in subsequent sessions. For example, in the Targ-Puthoff series with the subject Price the original transcript contains the phrase 'I've been trying to picture it in my mind and where you went yesterday on the nature walk,' the nature walk being the actual title used in the experiment to identify the previous day's target. This tells the judge not to match this particular transcript with the nature walk target. Thus statements of this kind need to be edited out. So when a feedback procedure is used the policy of no editing is unacceptable. Yet it is just in the case where feedback is used that the danger of bias from editing is greatest. Suppose that, subsequently to feedback, the percipient says 'What I really meant was this, not this' and asks for a correction to be made to the transcript to clarify his meaning. This request has to be firmly refused. But the percipient is, after all, the 'author' of the transcript and such a refusal may easily appear unreasonable. Thus the temptation is great to let rigorous standards slip. Suppose, then, that in view of these difficulties a policy of no feedback is adopted. Then we have somehow to persuade the agent and percipient not to communicate for the two weeks or so of the series. Worse, since one of the experimenters will always be present with the percipient during sessions and may cue him, whether intentionally or not, it is desirable that this experimenter should also remain ignorant of the target choices until the end of the series. Only an exceptionally disciplined experimenter is likely to submit to this constraint.

A number of remote viewing experiments have been conducted on a 'precognitive' basis: the selection and viewing of the targets has been carried out after completion of the recording of the percipient's impressions. Where the time interval is long enough, say one day or more, this allows enough time for the transcription and any necessary editing before target selection. Provided advantage is in fact taken of this opportunity this arrangement eliminates the possibility of introducing bias into the transcript. On the other hand it increases

the opportunity to introduce bias into the target selection or description; to avoid this the percipient's response material has to be kept secret for the whole period of a day or more. To ensure this without placing unacceptable restrictions on the percipient's freedom of action must present formidable problems. Yet no author of remote viewing reports seems to have mentioned these problems or discussed how they are handled in practice.

In some experiments the time order in which the targets were used has been communicated to the judge, either inadvertently by having the targets to be judged listed in chronological order of their use or simply by allowing the date of use to remain on the record of each target as sent for judging. When this happens the cues provided by allowing feedback may vitiate the experiment totally. In the Price series already mentioned, Marks and Kammann [9] showed that such cues were sufficient to explain the whole experimental effect: judges were able to piece together almost the whole sequence of targets and transcripts in their correct time order on the basis of such cues alone and without reference to the actual content.

Data selection bias and fraud

Just as the editing out of misses, whether in the target material or the transcripts, can increase the proportion of hits, so can the elimination of complete sessions which are unsuccessful. Usually a convincing pretext can be found for such an elimination: if it is the first session it was 'only a trial run'; if it is the last it is 'only an afterthought'; if it is in the middle perhaps the percipient said 'I don't feel I am having any success today', or conditions were different in some respect and the experimenter can persuade himself that this particular trial does not really belong with the series. In any case the experimenters have to satisfy only their own conscience: should they fail to mention the omission in the published report no one will ever know. Once again, the temptation to let standards drop is strong. The best safeguard is to deposit in advance with a reliable and independent party an unambiguous statement of the number of sessions to be included and the intended timing (e.g. 'twice a day for 5 consecutive days beginning 12 August'). No researcher appears to have ever taken this precaution. Some, however, do appear to have given conscientious attention to this issue. Dunne, Jahn and Nelson [5], for example, give details of cases which they class as 'informal' and make it possible for the reader to include or exclude them at will.

Finally, bias can be introduced by motivated selection of judgings. In the typical remote viewing experiment the material is judged and re-judged by a number of different judges, perhaps in order to smooth out the effect of eccentric judgments. Each judge will get a slightly different result and some results will be more favourable than others. If the less favourable judgings are simply dropped from mention in the report, a bias towards success can be created. Again, this does not necessarily imply a consciously dishonest experimenter: there are always reasons which can be found, after the event, for excluding a particular judging. The very fact that a certain judging is less favourable may even be seen, by an over-motivated experimenter, as evidence that it was performed carelessly and hence as justifying its rejection. Once again, such biases are best eliminated by preparing in advance a detailed plan of the research, including an exact identification of the judges who will be used. But again, this is a tiresome constraint and one which no researcher seems to have adopted.

No discussion of research into anomalous phenomena can avoid considering the question of fraud. The possibility of outright fraud by an experimenter can only be guarded against by demonstrating success under many different experimenters. But short of this extreme there is the possibility of fraud by the percipient. What are the options for a percipient wishing to obtain success by cheating?

If people can volunteer to serve as percipients or agents then a simple plan is for a pair to work in collusion in these two roles. They can agree in advance that the percipient is to stress some specific feature in each session (for example, circular shapes) and the agent can give special attention to such features in his commentary, sketch or photography. This plot could be foiled by the strategy already mentioned: requiring that all target material be frozen before the target selection is made. But if the pair are skilled magicians willing to go to some trouble it would be extremely hard to defeat them. Indeed a skilled magician as percipient hardly needs the cooperation of the agent: he can arrange for a confederate to follow the agent's car and telephone a description of the target area to another confederate near the laboratory. Innumerable methods are available for the latter to send information to the percipient in the laboratory during the session. If the percipient is informed eventually of the targets ('feedback') a simpler system is collusion between him and one or more of the judges. In some laboratories it seems to be common practice to select percipients and judges from among the same group, essentially the laboratory staff. To demonstrate success spread over several percipients would provide some safeguard; however, only a small number of 'free' hits are needed to yield a 'chance' expectation that will be computed as statistically significant using the standard procedure. (If n sessions are matched and 2 correct matches are guaranteed, the expectation for the sum of ranks is $2 + (n-2)(n-1)/2$ and this value appears as significant ($p < 0.05$) for all values of n in Table 1 in Morris [11].) In other words, in the judging situation most widely used only a small amount of trickery is needed to offer a better than even chance of significant success. Thus to eliminate the fraud hypothesis one needs to demonstrate success over quite a larger number of percipients.

Statistical evaluation

This leads, finally, to a consideration of the problem of statistical evaluation. There seems to be a surprising lack of concensus on the best method of evaluation among the different researchers in this field. However, at least all agree on one basic principle, that the judges should assess targets against percipients' descriptions in some kind of blind matching process in which descriptions intended for different targets are judged together with those intended for the actual target. Within this framework a variety of methods have been used. The most important distinction is perhaps between those methods in which the judging is confined to descriptions obtained in the same experimental series, typically a group of 5–10 sessions, and those in which material from other series is brought into the judging process. The latter system is used, for example, by Dunne et al. [5] in their extensive research and seems to have been more recently espoused by Puthoff. However, it is subject to serious criticism. The pool of 'control' material is built up over a period from experiments carried out in different places and at different times. The empirical probabilities of success are then applied to an experiment in which the time and approximate place are (in most cases) known to the percipient. But response preferences are likely to vary by region within a country and over time. (Perhaps in winter percipients are more likely to describe indoor scenes; near the sea or a lake they may more often mention water as part of the scene; perhaps bright colours of cars or clothes were more popular in the early 1980s than today; and so on.) It cannot be safely assumed that the probability of a hit established at another

time and place is a constant of nature and applies to an experiment carried out here and now.

Thus the only convincing method is to restrict the judging to a comparison of the material obtained during the different sessions of the same experimental series. Most remote viewing workers seem to have used this approach. Within this framework there is still a variety of statistical approaches possible. Rather than embarking on a lengthy discussion of these possibilities and their different limitations or defects I describe here the method which appears to me the least subject to criticism.

Let n be the number of sessions taken together for matching purposes – typically between 5 and 10. Then the basic requirement is to fill out an $n \times n$ matrix whose elements measure, in some way, the judged degree of resemblance between each description and each target. Provided judging is blind, the method to be described is valid no matter how these figures are obtained. For example, judges may be asked to rank each target from 1 to n against each description and these ranks may be entered directly in the matrix. (In this case a low score represents success.) Or judges may be required to rate the resemblances on some scale, perhaps taking separate account of each statement in the description. Whatever the method used, the matrix of these scores may be called the actual score matrix. The statistical analysis then proceeds by considering the set of all matrices obtained by rearrangement of the rows (or alternatively, of the columns) of the actual matrix in all possible $n!$ permutations. For each one of these permutations the sum of the scores in the diagonal is obtained. If the scores are such that a high figure represents success, we obtain P = the proportion of the $n!$ permutations which yield a diagonal sum greater than or equal to that of the actual matrix. If a low score represents success (as in the case of rankings) then P = the proportion with a diagonal sum less than or equal to that of the actual matrix. It is a simple matter to program a computer to sum the diagonal for each permutation and obtain P directly. This method has been used, for example, by Puthoff and Targ,[13] and by Schlitz and Gruber[16]. Note that, if several judges are used all judging the same material they cannot be treated as independent: their results should be totalled (or averaged in some other way) so that they all contribute to the same matrix.

The method described above overcomes all purely statistical problems in the evaluation of remote viewing experiments except for the problem of shared stimuli affecting agent and percipient together, and acting differentially between the sessions of one experimental series, such as the weather or the time of day. Note, however, that in so far as the method is incompletely defined (we have mentioned above that any method is permissible for representing the judgments numerically in the matrix) there is a loophole for a biased experimenter: he may try different methods and select only the one that gives the best score, failing to mention that he has done this. This argues for the use of a method that will be as transparent as possible. From this point of view one might recommend the ranking method: each judge ranks all targets from 1 to n against each description and the ranks awarded by the different judges are summed with equal weight. If all experimenters used this method there could no longer be any suspicions of motivated selection of the method to maximize scores.

Conclusion

A fair reaction to the above discussion would be that a properly conducted remote viewing experiment is a complex and delicate undertaking. Many precautions have to be taken, some of which are tiresome and inconvenient. There are numerous loopholes allowing bias to be introduced by an over-motivated and under-conscientious experimenter. Many of these are quite subtle and the experimenter will experience a strong temptation to ignore them. If the doubter is to be convinced by the evidence presented, the experimenter must make explicit that these precautions were in fact taken and that the loopholes were foreseen. Unfortunately, the sceptic who suspects an experimenter of lack of rigour is likely also to suspect his reports of lack of objectivity. The only safeguard against such criticism is a policy of full and open access to data. Some of the best known research in this field has fallen short of this standard (see correspondence in Nature: Marks, Kamman, Puthoff, Targ, Tart, Scott[7,8,10,14,17,18,20]). Further, scepticism is justifiably enhanced by the existence of numerous studies, seemingly conducted with essentially the standard methodology, which have obtained negative results (Allen et al.[1], Rauscher et al.[15], Karnes et al.[6]) to which may be added other negative outcomes that have not been reported in detail (Ornstein[12], Marks and Kammann[9], Butler[3]).

This article has discussed the methodological needs in some depth but the methodological practice in past experiments has been treated only superficially. Any thorough analysis of the existing body of evidence for remote viewing would have to go beyond the published reports, which are often brief and limited to a minimal description of procedures and results. One would need to inspect worksheets and interview participants. Even where the original experimenter is co-operative such work is time-consuming and frustrating. Moreover the social setting in which scientific work is conducted today is discouraging for the researcher seeking only to check the work of others: the contemporary ethos favours creativity, the breaking of new ground, the enlargement of the frontiers of knowledge. In this context it is difficult to find either the internal motivation or the necessary funding for purely critical studies. Thus, probing studies of the kind needed are rare, indeed in this field almost non-existent. Meanwhile, in the absence of hard criticism, shoddy work and pseudoscience are able to prosper, mimicking serious science by adopting its trappings while concealing the emptiness at the core.

Until a detailed appraisal going well beyond the published reports is available, the existing evidence for 'remote viewing' remains far from compelling. Certainly it does not justify the overthrow of well-established scientific principles.

1 Allen, S., Green, P., Rucker, K., Cohen, R., Goolsby, C., and Morris, R. L., A remote viewing study using a modified version of the SRI procedure, in: Research in Parapsychology 1975, pp. 46–48. Eds J. D. Morris, W. G. Roll and R. L. Morris. Scarecrow, Metuchen, N. J. 1976.

2 Bisaha, J. P., and Dunne, B. J., Multiple subject and long-distance precognitive remote viewing of geographic locations, in: Mind at Large. Eds C. T. Tart, H. E. Puthoff and R. Targ. Praeger, N.Y. 1979.

3 Butler, D. C., Personal communication cited by Karnes et al., ref. 6 below.

4 Dingwall, E., Abnormal Hypnotic Phenomena, vol. III. Churchill, London 1968.

5 Dunne, D. J., Jahn, R. G., and Nelson, R. D., Precognitive remote perception, Technical Note 83003, Princeton Engineering Anomalies Research Laboratory, Princeton University 1983.

6 Karnes, E. W., Ballou, J., Susman, E. P., and Swaraff, P., Remote viewing: failures to replicate with control comparisons. Psychol. Rep. *45* (1979) 963–973.

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

8 Marks, D., and Kammann, R., Information transmission in remote viewing experiments. Nature *274* (1978) 680–681.

9 Marks, D., and Kammann, R., The Psychology of the Psychic. Prometheus Books, Buffalo, N.Y. 1980.

10 Marks, D., and Scott, C., Remote viewing exposed. Nature *319* (1986) 444.

11 Morris, R. L., An exact method for evaluating preferentially matched free-response material, J. Am. Soc. psychical. Res. *46* (1972) 79–90.
12 Ornstein, R., Review of Mind-Reach, New York Times (March 13, 1977).
13 Puthoff, H. E., Targ, R., and May, E. C., Experimental psi research: implication for physics, in: The Role of Consciousness in the Physical World, pp. 37–86. Ed. R. G. Jahn. Westview, Boulder, CO 1981.
14 Puthoff, H., and Targ, R., Rebuttal of criticisms of remote viewing experiments. Nature *292* (1981) 388.
15 Rauscher, E. A., Weismann, G., Sarfatti, J., and Sirag, S.-P., Remote perception of natural scenes shielded against ordinary perception, in: Research in Parapsychology 1975, pp. 41–46. Eds J. D. Morris, W. G. Roll and R. L. Morris. Scarecrow, Metuchen, N. J. 1976.
16 Schlitz, M., and Gruber, E., Transcontinental remote viewing. J. Parapsychol. *44* (1980) 305–317.
17 Scott, C., No 'remote viewing'. Nature *298* (1982) 414.
18 Targ, R., and Puthoff, H., Information transmission under conditions of sensory shielding. Nature *251* (1974) 602–607.
19 Targ, R., and Puthoff, H. E., Mind-Reach: scientists look at psychic ability. Delacorte, New York 1977.
20 Tart, C. T., Puthoff, H. E., and Targ, R., Information transmission in remote viewing experiments. Nature *284* (1980) 191.

0014-4754/88/040322-05$1.50 + 0.20/0
© Birkhäuser Verlag Basel, 1988

The cold reading technique

D. L. Dutton

School of Fine Arts, University of Canterbury, Christchurch (New Zealand)

Summary. For many people, belief in the paranormal derives from personal experience of face-to-face interviews with astrologers, palm readers, aura and Tarot readers, and spirit mediums. These encounters typically involve cold reading, a process in which a reader makes calculated guesses about a client's background and problems and, depending on the reaction, elaborates a reading which seems to the client so uniquely appropriate that it carries with it the illusion of having been produced by paranormal means. The cold reading process is shown to depend initially on the Barnum effect, the tendency for people to embrace generalized personality descriptions as idiosyncratically their own. Psychological research into the Barnum effect is critically reviewed, and uses of the effect by a professional magician are described. This is followed by detailed analysis of the cold reading performances of a spirit medium. Future research should investigate the degree to which cold readers may have convinced themselves that they actually possess psychic or paranormal abilities.
Key words. Barnum effect; cold reading; psychic reading; astrology; paranormal reading; palmistry; spirit mediums.

The Barnum effect

That there is a sucker born every minute is the cynical slogan most often attributed to the great nineteenth century circus entrepreneur Phineas Taylor Barnum. Though there is in fact no record that he ever made such a remark, Barnum did claim that his success depended on providing in his shows 'a little something for everybody'. Both the cynicism and his recipe for success are relevant to understanding the persistent tendency for people to embrace fake personality descriptions as uniquely their own. This in turn gives a particular aptness to Paul Meehl's phrase, the Barnum effect, to describe the phenomenon [19].

Modern interest in the Barnum effect among psychologists dates from Forer's classic experiment in which a group of 39 undergraduate psychology students were given the Diagnostic Interest Blank [7]. A week later every student was provided with the same personality description, but was led to believe that each description was uniquely different, having been derived from the test results. The students were then asked to rate the accuracy of their 'individual' personality descriptions on a scale of 0 (poor) to 5 (perfect). Of the 39 students, only 5 rated it below 4, and no one rated it below 2 (average). The average rating was 4.3.

Forer's initial concern was with personality theory and assessment [8], and the relevance of this effect to understanding 'psychic' and apparently paranormal effects was not, to begin with, widely recognized. As late as 1962, P. A. Marks and W. Seeman were calling for the term 'Barnum effect' to be restricted to test-derived clinical personality descriptions of patients which are so general (and trivial) that they apply to everybody [18]. But the fact that Forer had obtained his generalized personality description not from standard texts in personality theory but from a newsstand astrology book indicated that the Barnum effect might be significant in contexts far beyond the clinical.

Forer's original personality description is perfectly serviceable even today. It was presented as numbered sentences; here it is as a continuous paragraph: 'You have a great need for other people to like and admire you. You have a tendency to be critical of yourself. You have a great deal of unused capacity which you have not turned to your advantage. While you have some personality weaknesses, you are generally able to compensate for them. Your sexual adjustment has presented some problems for you. Disciplined and self-controlled outside, you tend to be worrisome and insecure inside. At times you have serious doubts as to whether you have made the right decision or done the right thing. You prefer a certain amount of change and variety and become dissatisfied when hemmed in by restrictions and limitations. You pride yourself as an independent thinker and do not accept others' statements without satisfactory proof. You have found it unwise to be too frank in revealing yourself to others. At times you are extroverted, affable, sociable, while at other times you are introverted, wary, reserved. Some of your aspirations tend to be pretty unrealistic. Security is one of your major goals in life.' Parts of Forer's phraseology – 'sexual adjustment' – might be effectively put in more contemporary language, and there are current preoccupations which could be used to make the description even more immediately appealing to contemporary readers. But this is a description designed to be given to young people of both sexes, and it is remarkable how well it wears after almost 40 years. Slight variations on the original Forer description have been used in numerous studies in the 1950s, 1960s, and 1970s, all in one form or another replicating the remarkable susceptibility of subjects to the Barnum effect [3, 4, 6, 14].

There has been less agreement, however, on exactly why the Barnum effect works as well as it does. The most charitable interpretation would be to say that Barnum descriptions