

Why Agricultural Experts Disagree

James Shanteau

Kansas State University

The author can be contacted at:

Department of Psychology

Bluemont Hall 492

Kansas State University

Manhattan, KS 66506

(Ph: 785/532-0618)

(Fax: 785/532-5401)

(E-mail: shanteau@ksu.edu)

Running head: EXPERT DISAGREEMENT

Chapter to appear in: R. Saal (Ed.), *Farm Management Decisions with Climatic Risk* in

Toowoomba, Australia: DPI. (Revised 8/01)

Abstract

Why is it that agricultural experts so often seem to disagree? Although analysts often view such disagreements as a reason to mistrust experts, this paper argues instead that such should not be surprising. Such differences have been overlooked by previous investigators because of lack of understanding about the domain differences. That is, experts in some fields are more likely to disagree than in others. The purposes of this paper are: (1) to review evidence on disagreement by experts in various fields, (2) to outline a commonly-accepted conceptualization that questions the competence of experts because of these disagreements, (3) to explore domain differences in the degree of agreement between experts, (4) to suggest an alternate conceptualization of expertise that views disagreements as inevitable, (5) to look at implications for agriculture as well as for future directions for research on expertise.

Why Agricultural Experts Disagree

Analysts often seem perplexed when they observe sizable and consistent disagreements between professionals. This is particularly true when the professionals are considered experts. As shown by the following quotations, however, disagreements have been viewed historically as necessary.

“By different methods different men excel” (Churchill, 1764)

“The history of scholarship is a record of disagreements” (Hughes, 1936)

“The tough-minded . . . respect differences” (Benedict, 1940)

The position taken here is that disagreement between experts is not surprising, but rather is tied to particular domains of expertise.

The purpose of this paper is to explore the connection the extent to which experts agree or disagree and their area of specialty. The paper is organized into five sections as follows. First, there is a review of the literature on disagreement between experts. Second, a commonly held view that experts must agree or be considered incompetent is discussed. Third, the role that domain differences play in disagreement between expertise is considered. Fourth, an alternate perspective is offered that posits that experts in many domains are likely to disagree. Finally, the paper concludes with implications for agricultural and future research directions.

Background

Since the start of systematic analyses of decision making in the 1950's, investigators have expressed surprise and dismay at the extent to which experts disagree. For example, if we ask two agricultural experts to assess the viability of a farm, the expectation of most analysts is that they should come to the same conclusion. If they arrive at different conclusions, then we wonder whether they are as skilled as they claim.

In a seminal paper, Einhorn (1974) argued that *consensus* or between-expert reliability is a necessary condition for expertise. He found, however, significant differences in diagnoses by three expert medical pathologists. The average between-expert correlation (r) was .55 (where .0 is chance and 1.0 is perfect). In comparison, weather forecasters have been widely reported to have high consensus values, $r = .95$, for short-term predictions. Thus, there is considerable agreement between different weather forecasters (Stewart, Roebber, & Bosart, 1997).

It is also possible to examine *internal consistency*, the extent to which a single expert says the same thing in similar situations. For pathologists, Einhorn (1974) reported a within-expert consistency r of .50. For weather forecasters, on the other hand, the internal consistency is near perfect, $r = .98$ (Stewart, Roebber, & Bosart, 1997).

In a study of four professional livestock judges, experts were asked to evaluate overall breeding quality of swine (Phelps, 1977). Despite a high level of internal consistency (average $r = .96$), the consensus agreement was much lower, $r = .50$. Apparently, livestock experts have internally consistent strategies, but they do not agree with each other about what those strategies should be.

Comparable results have been reported for other types of agricultural judgments. For instance, grain inspectors were found to have a consensus value between judges of $r = .60$, with internal consistency of $r = .62$ (Trumbo, Adams, Milner & Schipper, 1962).

In non-agricultural domains, the values are often lower. For example, Hoffman, Slovic, and Rorer (1968) and Goldberg and Werts (1966) reported consensus values of less than .40 for judgments by professional stockbrokers and clinical psychologists. The internal consistency values were slightly higher with correlations of just over .40.

Several studies have explored whether between-expert consensus increases with experience. Ettenson, Shanteau, and Krogstad (1987) found that between-auditor correlations increased from .66 to .76

to .83, for students, audit seniors (mid-level), and full partners, respectively. Messier (1983) reported similar results – audit partners with more than 15 years experience had greater consensus than partners with less experience.

These results suggest three conclusions: First, experts in a variety of domains often disagree; the consensus correlations range from .40 to .60. Second, agricultural experts such as weather forecasters show higher levels of agreement, with r values up to .95. Finally, for nearly all domains, the internal consistency values are higher than the between-expert consensus values.

Experts-Should-Agree Argument

The less-than-impressive consensus correlations in most studies of expertise led many analysts to question the abilities of experts in general. Following Einhorn' s logic, these investigators assumed that agreement is a necessary condition for expertise. The lack of agreement, therefore, suggested that “e x-perts are no damn good” (Gettys, personal communication, 1980). This interpretation of reliability data apparently derived from an implicit five-part argument about experts:

(1) For tasks performed by experts, there is assumed to be a “gold standard” or unique “ground truth.” If this truth is readily accessible, anyone can obtain it directly. For expert tasks, however, the truth is outside the realm of common knowledge of most people That is why we need the experts.

(2) Because of their special skills and experience, experts should be able to tell us about this “ground truth.” That is, experts can access what others cannot access.

(3) Since by definition there can be only one “ground truth,” all experts should give us a single correct answer. The special abilities of experts thus allow them to make the same decision.

(4) If experts disagree, then someone is wrong – they cannot all be correct. Some (or all) of them are not real experts. Thus, disagreement is a reflection of incompetence.

(5) Since lay people do not know which of the so-called “experts” are correct, the only safe course of action is to distrust all (or most) of them. Thus, disagreement between experts implies that we should be suspicious of their claimed special abilities.

This argument, of course, is not a formal chain of logic. However, it is implicit in the way that many analysts have interpreted the meaning of disagreements between domain experts.

Domain Differences

It is common knowledge that experts in different domains perform different tasks. Yet, decision analysts persist in treating all experts alike, so that the term “expert” is used generically. For instance, Kahneman (1991) concluded, “there is much evidence that experts (in general) are not immune to the cognitive illusions that affect other people.” This may be true in some domains. But in some domains, such as weather forecasters, experts show little sign of biases or “cognitive illusions.” Thus, despite the generalizations drawn about experts, there are many exceptions to the rule.

In an effort to account for these domain differences, I constructed Table 1 to differentiate between those domains where experts do well and those where experts do not. The table is based on a continuum from high to low competence (see Shanteau 1992a,b for earlier versions of this table). In the left column are those domains where experts make aided decisions using Decision Support Systems (DSS) or other computerized tools, e.g., in weather forecasting. The next column contains domains where experts make skilled, but largely unaided decisions, e.g., livestock judges. The third column lists domains where experts show limited competence, e.g., clinical psychologists. In the last column, the behavior of experts in the last column is close to random, e.g., stockbrokers.

It should be noted that assignment of domains within the table was based on my review of the literature. That is, the level of competence is derived from the assessments of researchers who studied each domain.

=====
Insert Table 1 About Here
=====

There are many ways to describe the differences in this table (see Shanteau, 1992a,b). For present purposes, it is sufficient to observe that domains to the left (more competent) side involve stable (*static*) properties. That is, the stimuli and the problem “hold still” for experts to evaluate. The domains to the right (less competent) side involve more changeable (*dynamic*) properties. Thus, the stimuli and problem are less stable, harder to specify – more like “moving targets.”

Another way to view this distinction is to note that most domains to the left side of the table involve physical or natural properties, whereas the domains to right involve human behavior. This may reflect the fact that scientific advances have a longer history of development in left-side domains. In contrast, the domains to the right have relatively young scientific histories. Given this difference, it makes sense that expert competence will be higher to the left side and lower to the right side.

The question remains whether this distinction between domains is predictive of different levels of agreement. To examine this proposition, Table 2 shows consensus values for a variety of domains. Two specialties are listed under each category, with the between-expert agreement (consensus) given as average correlations. As can be seen, the mean consensus r value for weather forecasters is .95, whereas the average value for livestock judges is .50. The consensus values for clinical psychologists, and stock forecasters are .40, and .32, respectively. Comparable results for other domains appear in the second line. The trend supports the prediction outlined above – better-structured domains lead to high consensus values and less structured social-science domains lead to less consensus.

=====
Insert Table 2 About Here
=====

For comparison, the average within-expert reliability (consistency) correlations for these same domains are listed in Table 3. The trends are similar, with better-structured domains leading to higher internal consistency. As expected, the consistency values (except for pathologists) are higher than the corresponding consensus values in Table 2. In two domains (livestock judges and polygraphers), there are notable discrepancies between consensus and consistency correlations. These apparently reflect the existence of different ‘schools of thought’ as to how experts should do their jobs in these fields.

=====
Insert Table 3 About Here
=====

Analysts’ View of Experts

Investigators in artificial intelligence, expert system design, cognitive science, systems analysis, and computer science have all concluded that experts are superior decision makers. That is why knowledge engineers build computer simulations around what experts know. Similarly, most domain-specific researchers (such as in medicine and weather forecasting) view experts as possessing unique information essential for making good decisions. In short, investigators in these disciplines see human expertise as something to be emulated.

In contrast, decision analysts have concluded that experts are flawed and prone to making simple errors (e.g., Kahneman, 1991). Moreover, experts and novices are viewed as sharing the same shortcomings. For instance, Tversky (quoted in Gardner, 1985, p. 360) stated, “whenever there is a simple error that most laymen fall for, there is always a slightly more sophisticated version of the same problem that experts fall for.”

Decision analysts have apparently overlooked the fact that there are important domain differences between areas of expertise. For instance, both weather forecasters and livestock judges are skilled pro-

professionals. Yet, there is a major difference for these two fields in the level of disagreement between experts. The former is based on a well-developed science, whereas the latter is based more on “informed judgment.” It should not be surprising, therefore, to find livestock experts disagreeing more with each other about their judgments than weather forecasters.

Importance of Domains

The position taken here is that previous analysts have unknowingly adopted the *experts-should-agree* view of expertise. As argued above, this view implies that disagreement between experts is a sign that something is wrong. Moreover, that leads to the conclusion that experts are not as skilled or as competent as they claim to be. In this section, I will propose an alternative perspective based on a domain-sensitive view of expert performance. This perspective is based on a five-part argument:

(1) The primary job of an expert is not to make decisions but to help clients reach a broadly defined target state. For example, the goal may be to help policy makers design better land-management strategies or to increase the profit potential of a farmer. These goals do not involve single answers, but instead require something more elaborate from the expert, such as a strategic plan.

(2) To reach the client's goal state requires dealing with multiple, constantly changing, dynamic factors. As noted by Klein, et al (1993), the situations faced by experts are different and more complex than the simplified situations studied considered by most analysts. Thus, experts work on problems that are much more complex than those studied in idealized settings.

(3) Using their knowledge and experience, the role of the expert is to recognize patterns and find consistencies in a dynamic problem space. The expert's job is to clarify the issues for the client. In other words, the challenge for an expert is “to make sense out of chaos.”

(4) Based on their experience and insights into the nature of problems, experts try to help clients clarify their thinking. Typically, an expert will often identify various alternate paths to the desired goal

states. The expert's role is to lay out the options and the consequences in a clear and comprehensible fashion for others.

(5) In the end, it is the client, not the expert, who actually makes most decisions. The expert offers insights and observations, but it is up to others to make and implements the final choice(s). Thus, the final responsibility for the decision rests on the client, not the expert.

The view is nicely summarized by the management consultant Golde (1970): "We seem to expect too much and the wrong things of our experts." That is, experts generally act more like knowledgeable consultants. Rarely do they function as the "all-knowing, single-answer decision makers" envisioned by many analysts.

When experts disagree, therefore, it is because they often see alternative paths to the client's goal state. In turn, savvy clients may seek out the views of various experts precisely because they want different perspectives on their problems. Thus, disagreements between experts are not only expected, but may actually be useful to consumers.

Implications for Research

By relying on a false belief in consensus, many analysts have unknowingly adopted a distorted view that leads them to expect that experts should always agree. The next section looks further at the research implications arising from this distortion.

Economic/Statistical Thinking

By drawing a parallel to economic/statistical theory, decision analysts have adopted a "single correct answer" approach to assessing expertise. In most quantitative assessments, we expect to find one and only answer to a question. When an expert (or anyone else) gives an answer different from the "correct answer," he/she is said to have a "bias" (Tversky & Kahneman, 1974). Moreover, when two or more experts give different answers, their claim of special competence is questioned (Einhorn, 1974).

The position here is that these analysts have relied on an inappropriate view of how experts function. For instance, the environment in which experts work is much different from that reflected in the idealized world envisioned by these analysts. The complex, changeable environment that experts actually operate in is considerably more complicated. In reality, problems rarely are simple enough to lead to single correct answers. Instead, there are multiple answers (or at least multiple routes to answers). If so, it should not be surprising to find that experts, especially in domains involving human behavior, often take different approaches to finding answers.

The underlying problem is that analysts misunderstand what experts do and what is expected of them. These investigators seem to think that experts see the world as they do – with simplifying assumptions and single-answer solutions. However, experts generally have a different worldview – with many complexities and contingencies, but with few optimal solutions. In addition, experts have a flexible approach to adaptation and are better at managing uncertainty (Shanteau, 1992a).

From the present perspective, therefore, disagreements between experts are expected. Although analysts view disagreements as evidence of incompetence, my view is that experts see disagreements as a more-or-less inevitable part of their job.

Implications for Agriculture

As reviewed by Shanteau (in press), there has been a long history of interaction between agriculture and psychology. Although often unappreciated, these interactions go back over a century. For instance, Hughes (1917) study of crop judges is one of the first studies of experts in any domain. His analysis of crop judges showed that corn rated highest by experts did not produce the highest yield. Wallace (1923) (later to become vice-president under Franklin Roosevelt) reanalyzed Hughes' data using path analysis. He found (1) corn judges generally with each other, but (2) their ratings correlated only slightly with crop yields.

In a later study of crop judges, Trumbo, Adams, Milner, and Schipper (1962) asked licensed grain inspectors to judge samples of wheat. Their analyses revealed that nearly one-third of the samples were misgraded. When graded a second time (a check of consistency), over one-third were given a different grade. They also found that more years of experience increased confidence, but not lead to higher accuracy. In fact, more experienced judges tended to overgrade the wheat samples (perhaps the original form of ‘grade inflation’).

One explanation for these errors in agricultural judgment was offered by Gaeth and Shanteau (1984). They noted that irrelevant materials (e.g., excessive moisture) significantly impacted the decisions of soil judges. They also found that training to compensate for irrelevancies was successful in reducing the impact of these irrelevant materials. Another approach to improving expert judgment was developed in weather forecasting; Murphy and Winkler (1977) found that precipitation forecasts could be improved using a feedback system based on Brier scores (a quadratic lost function). Since the introduction of Brier scores, the accuracy of weather forecasts has increased dramatically (Stewart, et al., 1997).

Thus it appears that agricultural experts perform reasonably well, although there is clearly room for improvement, e.g., they are not always internally consistent. Moreover, there are sizable domain differences, i.e., there is greater agreement between experts in some areas than others. Experts working in domains with better-defined properties, such as weather forecasting, show sizable agreement. In contrast, experts in domains with more subjective standards, such as grain judging, perform less well.

Dawes (personal communication, 1987) offered an insightful observation about an earlier version of Table 1. The performance standards expected by clients are different for the left and right sides. Weather forecasters are allowed to make occasional mistakes. However, managers and stockbrokers

are expected to be correct almost all the time. That is, in the less predictable (right-side) domains, experts are held to higher standards of performance.

This is important for agriculture, since with increased mechanization and computerization, there has been a shift in skills needed on farms. Traditionally, a farmer needed to be a jack-of-all-trades, with general skills many areas. Today, modern agribusiness places more demand on cognitive (thinking) abilities. For example, with the trend away from family farms to large corporate entities, there is a greater need for farmers with management abilities. Yet, these are precisely the psychological skills associated with right side of the table (as opposed to the more technical skills on the left side). Whether farmers make such management decisions themselves or rely on expert consultants, the present paper suggests that these judgments will be lacking in competence and between-expert agreement.

Final Comments

As argued here, disagreements are a natural product of the various domains that experts work in. If disagreement between experts is not a focal issue, then what should be? Let me suggest three useful goals for future research on expertise. First, as argued above the superiority of experts depends on their ability to distinguish between relevant and irrelevant information (Ettenson, Shanteau, & Krogstad, 1987; Jacavone & Dostal, 1992; Mosier, 1997; Schwartz & Griffin, 1986; Shanteau, 1992b). One goal in future research should be to learn how experts make these discriminations and to find ways to enhance the process, especially for right-side domains.

A second goal should be to understand the kinds of intellectual and physical tools used by experts to enhance their judgments. Experts seldom, if ever, make unaided judgments of the sort emphasized in laboratory research. In fact, analysts make use of the very tools denied their subjects. "The experimenters themselves, using tools and expertise, are able to perform (laboratory) tasks rather well" (Edwards,

1983, p. 511). The type of tools used by experts needs to be better understood. Perhaps it will be possible to borrow some of the tools used by left-side experts to assist those on the right side.

The final goal should be to develop insights into domain differences. As argued by Edwards (1983, p. 512), “we have no choice but to develop a taxonomy of intellectual tasks themselves. Only with the aid of such a taxonomy can we think with reasonable sophistication about how to identify among the myriad types of experts and the myriad types of tasks . . . just exact what kinds of people and tasks deserve our attention.” The analyses in Tables 1, 2, and 3 offer an initial start in building a taxonomy.

Research on such goals will help broaden our understanding of expertise. In contrast, concern about the supposed incompetence of experts based on disagreements offers little opportunity for expanding our insights about expertise. Instead, we should focus our efforts on analyses of relevance/irrelevance, tool usage, and domain differences.

References

- Ashton, R. H. (1974). An experimental study of internal control judgments. *Journal of Accounting Research, 12*, 143-157.
- Boring, E. G. (1927). *A history of experimental psychology*. NY: Appleton-Century-Crofts.
- Edwards, W. (1983). Human cognitive capacities, representativeness, and ground rules for research. In Humphreys, P., Svenson, O., & Vari, A. (Eds.), *Analyzing and aiding decision processes*. Budapest: Akademiai Kiado.
- Einhorn, J. (1974). Expert judgment: Some necessary conditions and an example. *Journal of Applied Psychology, 59*, 562-571.
- Ettenson, R., Shanteau, J., & Krogstad, J. (1987). Expert judgment: Is more information better? *Psychological Reports, 60*, 227-238.
- Gaeth, G. J., & Shanteau, J. (1984). Reducing the influence of irrelevant information on experienced decision makers. *Organizational Behavior and Human Performance, 33*, 263-282.
- Goldberg, L. R., & Werts, C. E. (1966). The reliability of clinicians' judgments: A multitrait-multimethod approach. *Journal of Clinical Psychology, 30*, 199-206.
- Gardner, H. (1985). *The mind's new science: The history of the cognitive revolution*. NY: Basic Books.
- Gigerenzer, G. (1993). The superego, the ego, and the id in statistical reasoning, in Keren, G. & Lewis, C. (Eds.), *A handbook for data analysis in the behavioral sciences: Methodological issues*. Hillsdale, NJ: Erlbaum.
- Golde, R. A. (1969). *Can you be sure of your experts?* NY: Award Books.

Hamilton, R. E., & Wright, W. F. (1982). Internal control judgments and effects of experience: Replications and extensions. *Journal of Accounting Research*, *20*, 756-765.

Hoffman, P., Slovic, P., & Rorer, L. (1968). An analysis of variance model for the assessment of configural cue utilization in clinical judgment. *Psychological Bulletin*, *69*, 338-349.

Hughes, H. D. (1917). An interesting corn seed experiment. *The Iowa Agriculturalist*, *17*, 424-425.

Jacavone, J., & Dostal, M. (1992). A descriptive study of nursing judgment in assessment and management of cardiac pain. *Advances in Nursing Science*, *15*, 54-63.

Kahneman, D. (1991). Judgment and decision making: A personal view. *Psychological Science*, *2*, 142-145.

Kida, T. (1980). An investigation into auditors' continuity and related qualification judgments. *Journal of Accounting Research*, *8*, 506-523.

Klein, G. A., Orasanu, J., Calderwood, R., & Zsombok, C. E. (1993). *Decision making in action: Models and methods*. Norwood, NJ: Ablex Publishing Corp.

Lykken, D. T. (1979). The detection of deception. *Psychological Bulletin*, *80*, 47-53.

Messier, W. F. (1983). The effect of experience and firm type on materiality/disclosure judgments. *Journal of Accounting Research*, *21*, 611-618.

Mosier, K. L. (1997). Myths of expert decision making and automated decision aids. In C. Zsombok & G. Klein (Eds.), *Naturalistic decision making*. Hillsdale, NJ: Erlbaum.

Murphy, A. H., & Winkler, R. L. (1977). Can weather forecasters formulate reliable forecasts of precipitation and temperature? *National Weather Digest*, *2*, 2-9.

Phelps, R. H. (1977). *Expert livestock judgment: A descriptive analysis of the development of expertise*. Unpublished doctoral dissertation. Manhattan: Kansas State University.

Phelps, R. H., & Shanteau, J. (1978). Livestock judges: How much information can an expert use? *Organizational Behavior and Human Performance*, *21*, 209-219.

Raskin, D. C., & Podlesny, J. A. (1979). Truth and deception: A reply to Lykken. *Psychological Bulletin*, *86*, 54-59.

Shanteau, J. (1989). Cognitive heuristics and biases in behavioral auditing: Review, comments, and observations. *Accounting, Organizations, and Society*, *14*, 165-177.

Shanteau, J. (1992a). Competence in experts: The role of task characteristics. *Organizational Behavior and Human Decision Processes*, *53*, 252-266.

Shanteau, J. (1992b). How much information does an expert use? Is it relevant? *Acta Psychologica*, *81*, 75-86.

Shanteau, J. (1999). Decision making by experts: The GNAHM effect. In Shanteau, J., Mellers, B., & Schum, D. (Eds.), *Decision science and technology: Reflections on the contributions of Ward Edwards*. Norwell, MA: Kluwer Academic Publishers.

Shanteau, J. (in press). Agricultural psychology. In W. E. Craighead & C. B. Nemeroff (Eds.), *Encyclopedia of Psychology and Neuroscience*. N.Y.: Wiley.

Slovic, P. (1969). Analyzing the expert judge: A descriptive study of a stockbroker's decision processes. *Journal of Applied Psychology*, *53*, 255-263.

Stewart, T. R., Roebber, P. J., & Bosart, L. F. (1997). The importance of the task in analyzing expert judgment. *Organizational Behavior and Human Decision Processes*, *69*, 205-219.

Trumbo, D., Adams, C., Milner, M., & Schipper, L. (1962). Reliability and accuracy in the inspection of hard red winter wheat. *Cereal Science Today*, *7*, 62-71.

Tversky, A., & Kahneman, D. (1974). Judgment under uncertainty: Heuristics and biases. *Science*, *185*, 1124-1131.

Wallace, H. A. (1923). What is in the corn judge's mind? *Journal of the American Society of Agronomy*, 15, 300-304.

Yates, J. F. (1990). *Judgment and decision making*. Englewood Cliffs, NJ: Prentice-Hall.

Author Note

Preparation of this manuscript was supported, in part, by *National Science Foundation Grant DMI 96-12126* and by support from *Federal Aviation Administration Grant 98-G-026*. The author wishes to thank Ward Edwards, Julia Pounds, and Rob Ranyard for their helpful comments on earlier versions of the manuscript.

Correspondence concerning this paper should be addressed to *James Shanteau, Department of Psychology, Bluemont Hall 492, Kansas State University, Manhattan, KS 66506*.

Table 1

Progression of Domains from High to Low Performance

| Highest Levels of Performance..... | | Lowest Levels of Performance | |
|------------------------------------|--------------------|------------------------------|-----------------|
| <i>Aided Decisions</i> | <i>Competent</i> | <i>Restricted</i> | <i>Random</i> |
| _____ | _____ | _____ | _____ |
| Weather Forecasters | Livestock Judges | Clinical Psychologists | Stockbrokers |
| Auditors | Grain Inspectors | Pathologists | Polygraphers |
| Test Pilots | Chess Masters | Parole Officers | Office Managers |
| Insurance Analysts | Photo Interpreters | Student Admissions | Parole Officers |
| Physicists | Soil Judges | Intelligence Analysts | Court Judges |

Note: Each domain was classified based on analyses of expert performance reported in the literature.

Table 2

Reliability (Consensus) Values for Experts

| | | | |
|------------------------------------|------------------|-----------------------------------|----------------|
| Highest Levels of Performance..... | |Lowest Levels of Performance | |
| <i>Aided Decisions</i> | <i>Competent</i> | <i>Restricted</i> | <i>Random</i> |
| _____ | _____ | _____ | _____ |
| Weather Forecasters | Livestock Judges | Clinical Psychologists | Stockbrokers |
| <i>r</i> = .95 | <i>r</i> = .50 | <i>r</i> = .40 | <i>r</i> = .32 |
| | | | |
| Auditors | Grain Inspectors | Pathologists | Polygraphers |
| <i>r</i> = .76 | <i>r</i> = .60 | <i>r</i> = .55 | <i>r</i> = .33 |

Note: The values cited in this table were drawn from the following studies (from left to right): Stewart, Roebber & Bosart (1997); Phelps & Shanteau (1978); Goldberg & Werts (1966); Slovic (1969); Kida (1980); Trumbo, Adams, Milner & Schipper (1962); Einhorn (1974); and Lykken (1979).

Table 3Reliability (Internal Consistency) Values

Highest Levels of Performance.....Lowest Levels of Performance

| <i>Aided Decisions</i> | <i>Competent</i> | <i>Restricted</i> | <i>Random</i> |
|----------------------------------|-------------------------------|-------------------------------------|----------------------------|
| _____ | _____ | _____ | _____ |
| Weather Forecasters $r = .98$ | Livestock Judges $r = .96$ | Clinical Psychologists $r = .44$ | Stockbrokers $r = <.40$ |
| Auditors $r = .90$ | Grain Inspectors $r = .62$ | Pathologists $r = .50$ | Polygraphers $r = .91$ |

Note: The values cited in this table were drawn from the following studies (from left to right): Stewart, Roebber & Bosart (1997); Phelps & Shanteau (1978); Goldberg & Werts (1966); Slovic (1969); Kida (1980); Trumbo, Adams, Milner & Schipper (1962); Einhorn (1974); and Raskin & Podlesny (1979).

Biographical Sketch

James Shanteau is Professor of Psychology at Kansas State University, where he served as Director of the Institute for Social and Behavioral Research. He has held visiting positions at the National Science Foundation, as well as the Universities of Michigan, Oregon, Colorado, and Cornell. He is Co-Founder and Past President of the Society for Research on Judgment and Decision Making. He is a Fellow of both the American Psychological Association and the American Psychological Society. His research interests include decision making by domain experts, models of judgment and decision making, and quantitative models of behavior.