Japanese Psychological Research
1997, Volume 39, No. 1, 1–11

Invited Paper

The status of the concept of intelligence[1]

EARL HUNT
Department of Psychology, University of Washington, PO Box 351525, Seattle, WA 98195-1525, USA

Abstract: Psychometric studies have shown that "general intelligence" should be broken down into the ability to apply learned solutions to new problems (crystallized intelligence) and the ability to deal with novel intellectual problems (fluid intelligence). This distinction has been amplified upon by studies of individual differences in information processing. Crystallized intelligence depends on the problem-solving schema that people have acquired and upon their efficiency in accessing information in long-term memory. Fluid intelligence is associated with the ability to access and manage relatively large amounts of information in working memory. Measures of fluid and crystallized intelligence are important predictors of objectively measured workplace performance. Studies of actual and simulated workplaces have shown that this is largely due to differences in people's ability to manage information and the speed with which the details of a job can be grasped.

Key words: individual differences, information processing, intelligence, long-term memory, schema, working memory, workplace performance.

We humans are sure that we are the most intelligent species on Earth, though sometimes our behavior may suggest another point of view. It is clear that not all of us think in quite the same way, and some of us think a little bit better than others. I want to explore just what this means.

The modern study of intelligence began about 100 years ago, when Sir Francis Galton called attention to the fact of individual differences in a variety of cognitive tasks. Galton emphasized what we would today call low-level cognitive tasks, such as simple reaction time. Then Binet utterly ignored Galton's methods of measurement in order to begin an intelligence testing program that met an applied need, screening students in public schools. This split became a tradition. For more than half a century the test developers, the psychometricians, and the laboratory-based experimental psychologists virtually ignored each other. From time to time people said we should get together, because after all we were studying the same human mind (Cronbach, 1957), but little was done until the early 1970s, when my own laboratory, and somewhat later Robert Sternberg (1977) and Benton Underwood (Underwood, Boruch, & Malmi, 1978), published empirical work uniting the laboratory and the testing center. Shortly after, the journal *Intelligence* was founded. Today, its articles regularly combine cognitive psychology and psychometric approaches. Work has prospered for about a quarter of a century, so it is time to stop and see where we are.

[1] Text of an invited address to the Annual Convention of the Japanese Psychological Association.

The modern psychometric view

Since psychometric tests are the popular definition of intelligence, let us begin with them. The first question to ask is not "How are the tests constructed?" but rather "Do the tests measure anything that matters?" The correlation between intelligence test scores and academic performance is between .3 and .5, depending on whether or not you wish to discuss observed correlations or population estimates based upon extrapolations from the observed statistics (Hunt, 1995). Similar correlations are found between test performance and performance in military and industrial training programs (Earles & Ree, 1992). What is less widely realized is that the tests predict industrial as well as academic performance (Hunter, 1986), providing that you are looking at the cognitive and not the social aspects of how people do their jobs. The distinction makes a difference. Figure 1 presents the results of a study by the U.S. Army, which found that intelligence test scores predict the technical aspects of job performance quite well, but are poor predictors of motivation.

Every measuring device implies a theory of the thing being measured. The best-known psychometric theory is Spearman's notion of general intelligence (Spearman, 1927). Spearman believed that performance on any intellectual task was determined by a person's general intelligence, g, augmented by a variety of special intelligences that were unique to specific testing procedures. While some people still accept Spearman's theory (e.g., Herrnstein & Murray, 1994), many psychometricians believe that individual differences in mental abilities are too complex to be represented by a single dimension. This position, which today is vigorously championed by Howard Gardner (1983), was first put forward by Leon Thurstone (1938).

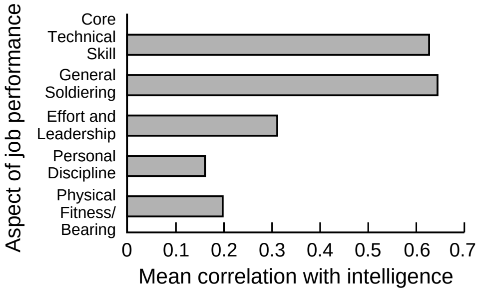

Figure 1. The estimated relationship between Armed Services Vocational Aptitude Battery (ASVAB) components and various dimensions job performance amongst U.S. Army enlisted personnel (McHenry, Hough, Toguam, Hanson, & Ashworth, 1990).

In the 1930s Thurstone developed a theory of intelligence in which people were described by their coordinates in an intellectual space whose dimensions were such things as verbal comprehension and arithmetic facility. Thurstone argued that rather than being just one thing, intelligence was composed of 8 to 12 separate abilities. Subsequently, more than 100 dimensions were proposed (Guilford, 1967). R. J. Sternberg (1990) has aptly called such theories of intelligence a geographic metaphor for the mind, albeit using a far more complicated geography than our trivial three-dimensional globe.

Now let us consider a third type of theory, *hierarchical* theories, such as those proposed by Horn, Cattell, and (somewhat earlier) Philip Vernon (Cattell, 1971; Horn, 1985; Horn & Noll, 1994). These theories are a compromise between Spearman's notion of general intelligence and the multiple intelligence ideas of Gardner, Guilford, and Thurstone. Horn and Cattell maintain that general intelligence consists of two correlated dimensions, fluid intelligence (Gf) and crystallized intelligence (Gc). In addition, they acknowledge a third dimension of ability, *spatial-visual reasoning* (Gv), which is statistically almost independent of Gf and Gc. Loosely, fluid intelligence is the ability to figure out ways of attacking novel problems, crystallized intelligence is the ability to apply previously learned solution methods to the problem you are currently facing, and spatial-visual reasoning is just what it says it is, the ability to reason about images and locations in space. There is more to the theory, but this brief summary is enough detail for our present purpose.

Using statistical techniques that were not available to the early test developers, John Carroll (1993) conducted a massive analysis of many of the most important data sets on which our theories are based. He compared the different theories using standard analytic methods on the same evidence. Carroll found that the hierarchical model was the best for virtually all the data. To repeat, the basic distinction in the "winning" hierarchical theory is between the ability to solve new problems, the ability to apply old solutions to current problems, and the ability to use visual imagery and spatial reasoning. I will assume this distinction for the remainder of this paper. However, I want to qualify Carroll's conclusions in an important way, by considering some data that were outside of his survey.

All psychometric theories carry with them the implicit assumption that the relation between different types of intelligence are the same at the top and the bottom of the intellectual spectrum. However, this assumption is not true. Spearman's general intelligence theory may be a good description of individual differences at the lower end of the intellectual range, while the differentiated Horn-Cattell model is appropriate at the upper ends.

The best evidence for this claim is a study by Detterman and Daniel (1989) of the sample used to validate the Wechsler Adult Intelligence Scale-revised (WAIS-R). Detterman and Daniel first split this large data set into subgroups consisting of the quintiles of scores on just one of the subtests. They then computed the mean correlation between each pair of *other* subtests within each quintile. If the assumption of unchanging factor structure is correct, the correlations should be lowest in the mid-quintile, and then should increase *symmetrically* as we move outward, to the quintiles either above or below the mid-quintile. The reason is that the range of scores in the outermost deciles is greater than the range of scores in the deciles in the mid-ranges. But this is not what happened. The data are displayed in Figure 2. As can clearly be seen, the change in correlations is not symmetrical. Subtest correlations are high at the low end and low at the high end of ability.

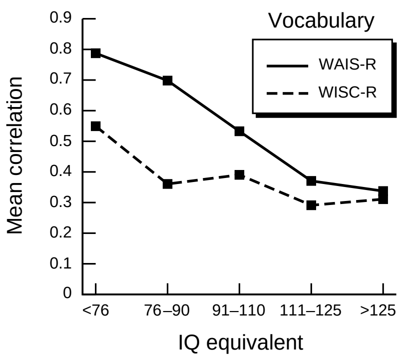

Figure 2. Average intercorrelation between subtests of the WAIS as a function deciles of scores of other subtests of the WAIS, which were removed from the correlation matrix. From Detterman and Daniel, 1989.

This finding is not unique to the WAIS. Figure 3 shows a similar analysis that Derek Chung and I made of the intelligence test scores available for the National Longitudinal Study of Youth, a large, carefully constructed population of the younger American workforce. The similarity between our results and Detterman and Daniel's is striking.

To summarize, psychometric theory began with the implicit assumption that there is a single, pervasive trait of general intelligence, and that it mattered both in academia and in the workplace. After about a century of debate, the following picture has emerged:

1. Intelligence is one of the most important factors determining academic and workplace performance. Clearly, no one factor determines success in human endeavors. However, intelligence cannot be ignored.
2. The concept of general intelligence has to be split into three separate general abilities; fluid, crystallized, and visual-spatial intelligence. Fluid and crystallized intelligence are at least moderately correlated.
3. While high-level abilities appear to be differentiated, low-level abilities, or lack of intelligence, is general.

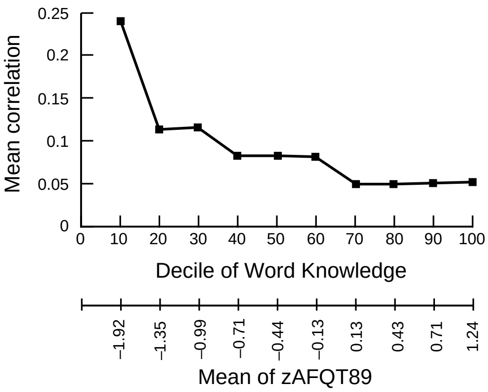

Figure 3. Average intercorrelation between subtests of the ASVAB, as a function of deciles of scores on the word knowledge subtest. Computations based on the National Labor Survey of Youth (NLSY) data bank. The computations were carried out for this article by Earl Hunt and Derek Chung.

A cognitive psychology view of individual difference

Now let us look at intelligence from the perspective of a cognitive psychologist. Theories in cognitive psychology begin with a commitment to some principles about how the mind processes information in general, without concern for the semantic content of the information being processed. Using an analogy with the design of computers, this will be called the *system architecture* level of cognitive theory. Most cognitive psychology theories today are based on the "blackboard" system architecture shown in Figure 4. Memory is divided into two broad classes: immediate memory and long-term memory. The role of immediate memory is to provide a workspace representing what the thinker is attending to at the moment. Long-term memory is the repository for information acquired in the past. Such information falls into two broad classes: declarative information, about facts, and procedural information, about how to do things. To illustrate, my knowledge that I was in Seattle yesterday morning is declarative information, while my knowledge of how to drive a car or solve a set of linear equations is procedural information. Procedural information is very important. According to the blackboard model, individual acts of thinking are driven by pattern recognition. Much of our memory consists of rules for behavior and the circumstances in which we should use them.

An architectural-level theory describes a potential for thought, just as a computer, as a physical device, establishes a potential for computation. Theories of thinking about something (e.g., theories of chess playing, language comprehension, or mathematical theorem proving) augment the architectural-level theory with a representational-level theory that specifies what knowledge the individual has and how that knowledge is used. Clever memorization can greatly augment the information processing capacity provided by our brains. A classic example is provided by short-term memory. Many years ago Miller (1956) noted that our memory for repeating letters is quite short. Given a random letter string we can repeat back at most eight or nine letters. However, if

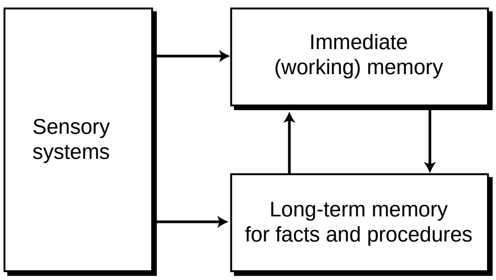

Figure 4. The Blackboard Model of human information processing. Working memory contains a representation of the problem before us. This representation is acted upon by declarative and procedural information in long-term memory.

the letters are organized into a code that is meaningful to us, as in:

JAPAN WINS OLYMPIC VOLLEYBALL
GAME FROM USA

we do much better. At a more complex level, school children learn how to apply the Pythagorean theorem to solve problems in geometry. Much later they learn to use linear regression analysis in algebra. These are *schema* that, once learned, permit us to organize our problem-solving efforts in ways that apply our mental architecture in an efficient manner. In order to understand a person's thinking you must understand the capabilities that that person has for processing information in general, which are provided by the system architecture, and you must understand the schema and chunking procedures that the individual uses to organize information about a particular topic. Where are the individual differences in this scheme?

During thinking our brains must ship information around from one center of neural activity to another. This does not mean that we have wires inside our head, analogous to wires inside a computer, but it does mean that the mind must contain channels of communication between different centers of mental action. Since the brain supports the mind, the brain must contain information-processing channels that move signals from one site of brain action to another. Jensen (1987) has argued that individual differences in neural channel capacity, speed, and reliability are important contributors to individual differences in intelligence, and that such differences can be revealed by studying the speed and accuracy with which people make very simple choices.

There is a good bit of evidence supporting this idea. The correlation between intelligence test scores and the time required to make simple choices, such as the *choice reaction time (CRT)* paradigms is about .3 (Jensen, 1987; Palmer, MacLeod, Hunt, & Davidson, 1985; Vernon, 1983; Vernon & Kantor, 1986). Very much the same point has been made in *inspection time (IT)* studies, in which people must make judgments about very rapidly presented visual displays. In general, people with high intelligence test scores extract information from visual displays more rapidly than people with low scores.

It is clear that one of the major individual differences in intelligence is the sheer speed with which people can make simple decisions. Interestingly, one of the best-documented findings in geriatrics is a progressive slowing of performance on virtually all information processing tasks as a function of aging. This slowing can be explained by a simple model that assumes that as we age our neural information-processing channels simply become less reliable (Cerella, 1990; Myerson, Hale, Wagstaff, Poon, & Smith, 1990). There is a temptation to think that this may account for the equally well documented age-related drop in fluid intelligence (Horn, 1985). While this may be so, further exploration is required. Many phenomena that are linked with age (e.g., loss of hair in males) have nothing to do with declines in intelligence.

The point about aging can be generalized. Information processing speed cannot explain all the relationships between information processing and intelligence, because the relationships increase as the information processing tasks become more complex (Vernon & Kantor, 1986). In fact, partialling out CRT scores does not markedly influence the relationship between reading and intelligence test measures (Palmer et al., 1985). A closer look at the relationship between intelligence and information processing measures is in order.

Thinking involves the manipulation of a mental representation held in immediate memory. Therefore we ought to find that measures of the performance of the immediate memory system are related to measures of complex thinking. Indeed, they are.

One of the best-known paradigms in information processing psychology is the short-term memory scanning task developed by S. Sternberg (1966). In this task people are shown a *memory set* of between one and six items, and then immediately asked if a probe item was a member of the memory set. The time to respond to the probe rises linearly with the number of times in the memory set, so the slope of the function (the *scanning rate*) is taken as a measure of speed of access to information in immediate memory. In one of our initial studies (Hunt, Frost, & Lunneborg, 1973) we found that college students who have high scores on mathematical aptitude tests tend to be rapid scanners. However, for a variety of reasons, we failed to follow up this finding. That was our misfortune. Subsequently, S. Sternberg (1975) reported that the slope varies systematically in widely defined populations, from 40 ms per item in college students to over 100 ms in elderly adults and some patient groups. Meanwhile we had found that scanning rates could be varied by the use of barbiturate drugs, which subjectively often leave people feeling dull (MacLeod, Dekaban, & Hunt, 1978). However, to my chagrin, we did not follow up this line of investigation.

Others did. Daneman and Carpenter (1980, 1983) developed a way of measuring a person's *reading span*, the extent to which people can hold extraneous information in immediate memory while processing sentences. They and their colleagues have since shown that reading span size is correlated with performance in a host of tasks that require fairly complicated linguistic reasoning. Patrick Kyllonen and Raymond Christal (1990) extended these results by showing that individual differences in complex reasoning tasks, such as the analysis of electronic circuits, can be almost entirely accounted for by individual differences in working memory. Taking these results together, it appears that reasoning tasks of the sort usually described as "requiring fluid intelligence" are almost exactly those that require fairly large immediate memory capacities.

Now let us shift our attention from parameters related to the active processing of information to parameters related to the process of extracting information from long-term memory. This process is essential for thought, because most human thought is guided by the retrieval of previously acquired information and procedures. A particularly apt example is reading, which is an essential cognitive skill in modern society. Rapid readers must make sight–meaning correspondences within a few milliseconds. In English, for instance, the good reader must associate the arbitrary letter sequence CAT with the concept of "small, furry, domestic animal." Lexical entries must be distinguished from nonlexical ones, such as CAK, which is not a word in English although it does obey English language orthographic and phonetic conventions. Recognition must be invariant over some scripts, but at the same time be sensitive to small changes in the visual stimulus. A, *a*, and a all refer to the same letter but b and d are different.

During the 1970s and 1980s my colleagues and I (Hunt et al., 1973; Hunt, Lunneborg, & Lewis, 1975; Hunt, Davidson, & Lansman, 1981; Palmer et al., 1985; see also Bell & Perfetti, 1994) showed that there is a moderate correlation (*r* in the .3 and .4 range) between the speed with which people can do a simple recognition task and the accuracy with which they can perform much more complicated linguistic tasks, such as paragraph comprehension. Memory counts, which is not surprising. But there is more, for these findings open the door to a consideration of the different ways in which memory counts.

In reading, the cues to word identification are closely linked to the physical nature of the stimuli. English readers recognize CAT, *CAT*, and, with slightly more difficulty, CAT, as permissible distortions of our prototype of the written word "cat." The ability to make direct recognitions of this sort is certainly important in thought. However, much of our more complicated thinking depends upon making an indirect connection between the present situation and past knowledge, by linking them both to an abstract situation. For instance, school children do not learn how to count dogs, cats and marbles. They learn arithmetic schema that let them do simple counting problems about anything (Kintsch & Greeno, 1985). The same principle applies to adult reasoning. Physicists use concepts like "balance of forces" and statisticians use concepts like "independent groups designs" as ready-made templates to guide problem solving (Larkin, 1983). Students learning physics do the same thing, except that their schema are neither as well organized nor as close to the laws of physics as their instructors might like (di Sessa, 1993; Hunt & Minstrell, 1994). The use of schema is by no means limited to mathematics and the sciences. Criminal judges determine their sentencing decisions by deciding whether the case at hand represents a situation where society needs protection, a criminal needs punishment, or an unfortunate individual needs help (Lawrence, 1986). The intelligent expert, magistrate or mathematician has developed procedures that facilitate transfer from a particular case into a schema, and then back out from a schema into a case-specific solution. Humans develop expertise in various fields by developing the schema that are appropriate to the life they live.

We acquire most of our schema by experience with specific situations. Therefore how a person codes his/her experiences will determine what is learned from them. If an experience is coded in terms of general principles, then the experience can be used to solve future problems, providing that the future problem is also coded in the same way. There are very wide individual differences in the facility with which people do this, a point that has been documented in the literature. In fact, some of the greatest failures in intelligent processing seem to be associated with people's failure to realize that they have information relevant to the task at hand. Rather than cite statistics, I will rely on an anecdotal observation I made in the course of studying how people learn scientific concepts in elementary physics and mathematics (Hunt & Minstrell, 1994).

Half a dozen science teachers participated in an exercise in which they measured the diameter and circumference of a number of circular objects (e.g., a tube) and then plotted the circumference as a function of the diameter. Not surprisingly they "discovered" the linear relationship:

$$\text{circumference} = A * \text{diameter}$$

and A was approximately 3.17. The teachers had "rediscovered" the value of π. Then my colleague Jim Minstrell posed the question shown in Figure 5. Suppose a steel band was placed around the Earth at the equator. The band would have to be 25,000 miles (40,000 km) long. How much longer would it have to be if the band were suspended 6 feet (2 m) above the equator? With one exception, the teachers asserted that the band would have to be very much longer. If they applied the schema that they had been using minutes before they would have found that the second band would be about 37 feet (12 m) longer than the first.

My purpose is not to poke fun at teachers. I suspect that many other professionals would have made the same mistake. The point of this story is that the transfer is difficult, because it requires that both the learning and the test situation be mapped into the same abstract representation. Forming the representation the first time is an exercise in problem solving/fluid intelligence, while recognizing that the representation can be applied is an exercise in crystallized intelligence. With this perspective we now return to psychometrics.

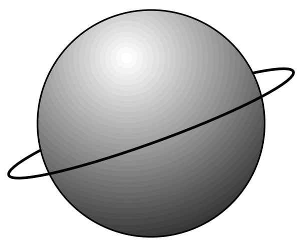

The equator is 40,000 km long. Suppose a steel band is placed 2 m above the equator. How much longer than the equator would the steel band be?

Figure 5. Question posed to science teachers shortly after they had plotted values of circular objects as a function of their radii. Most teachers failed to realize that the length of the band around the equator would be only slightly increased if it were raised 2 m above the earth.

Complementarities and importance

R. J. Sternberg (1990) has pointed out that when we explore a topic as complex as the mind we use metaphors to guide our thinking. Cognitive psychologists and psychometricians use different metaphors, but their ideas are beginning to converge. The picture that emerges from cognitive psychology studies of intelligence is strikingly complementary to Cattell and Horn's division of general intelligence into fluid and crystallized components. (The complementarity extends to visual reasoning, the Gv component, but there is not space here to discuss the evidence.) People faced with new situations (requiring Gf) must develop new representations of the current problems, thus imposing a heavy load on immediate (working) memory and upon the closely related ability to move information around in the brain/mind system. When problem solving depends upon previously acquired knowledge the information processing burden shifts to pattern recognition processes, which are a separate part of the human information processing system.

It follows that most learning is, in effect, problem solving rather than a rote memorization task. Therefore intelligence, as tested by conventional tests with large g or Gf components, should do a better job of predicting individual differences in performance during the acquisition phase of a task than during performance following acquisition. They do. The validity coefficients for intelligence tests are substantially higher when the criterion is performance in a workplace setting (Hartigan & Wigdor, 1989). This finding is sometimes interpreted as showing that intelligence tests predict only "school" performance. I disagree. I believe that the key issue is learning, not the fact that the learning is taking place in the school.

Figure 6 summarizes a U.S. military study of the performance rating of enlisted personnel as a joint function of their intelligence test (with class I-II personnel having the highest scores, and class IV the lowest) and the time that they had been in their military job. The relationship

Figure 6. The relationship between Armed Forces Qualification Test (AFQT) category, number of months on the job, and job performance (Wigdor & Green, 1991).

between job knowledge and intelligence becomes smaller the longer people are on the job, but it never completely goes away.

Much the same point can be made in laboratory studies, where learning is studied under more tightly controlled conditions, but over a much shorter timescale, usually days rather than months. Ackerman (1988) has conducted an extensive series of such experiments. He found that the correlation between intelligence test scores and performance on a simulated air traffic control test dropped, but did not disappear, as training was extended. In my own laboratory Susan Joslyn has conducted a study that extends Ackerman's results, using a somewhat more realistic simulator. The Raven Progressive Matrix, a test that shares no content with air traffic control, predicted the number of trials required to learn the controller's task to a set criterion ($r = .68$), but was a poor predictor of performance after training had been completed ($r = .07$).

Another implication of the cognitive psychology approach is that as people become more expert they will become more specialized. They will learn to develop the schema and exercise the cognitive skills that are relevant to their interests. This is consistent with Detterman and Daniel's finding (and our extension of it) that cognitive skills are more differentiated at the upper than at the lower end. However, the specialization argument also suggests that related skills should become more closely bound together. To use an extremely stretched metaphor, we are faced with something of a "big bang" theory of intelligence, where the g factor relates to time. At the low end, all abilities are closely bound together. At the high end, abilities form "galaxies," where related abilities are tied together into closely bound clusters, but the clusters themselves lie further and further apart.

Does this specialization actually happen? David Waller and I, in unpublished work in our laboratory, examined the correlations between subtests of the ASVAB in the NLSY sample. We estimated population correlations between subtest, using as samples either the highest or lowest deciles, as defined by removing one test and using it as a stratifying variable. As has already been shown (Figure 3, p. 4) the average correlations between subtests decrease when this is done. However, the highest correlations are virtually unchanged. For instance, two of the most highly correlated subtests in the NLSY data bank are the pairs (word knowledge, paragraph comprehension) and (word knowledge, general scientific knowledge). The correlations between these pairs are respectively .67, .68 in the 0th decile and .70, .70 in the 9th decile (averaging over deciles defined by different tests). On the other hand, the lowest correlation in the 0th decile is between clerical-perceptual skills and automechanics knowledge (.30). In the 9th decile this correlation is –.04.

These results are extremely important for industrial societies, albeit for not quite the reasons that some observers (notably Herrnstein & Murray, 1994) have suggested. Workforces are aging throughout the industrial world. At the same time, technological advances have accelerated changes in how the workplace operates (Hunt, 1995). We know that as adults age there are declines in fluid intelligence and slight increases in crystallized intelligence. This is of no matter in a static society, for what the worker sells to the employer is knowledge of how things have worked in the past. But if the future is not going to be like the past, the worker has to sell his or her ability to learn. This ability may be precisely what decreases with age. Viewed from the optimistic side, there is a fundable research opportunity here! We need to find ways to facilitate lifelong learning. If such research is not carried out, *and its results implemented*, older workers will be faced with economic insecurity at precisely the point in their lives at which they, as parents and responsible leaders, require maximum economic security so that they can meet the other demands of life. This effect is magnified, of course, by any tendency for women either to postpone parenting or to marry older men. Our society has yet to think through the profound interaction between technological change and the demographics of cognitive abilities.

About 40 years ago Cronbach (1957) urged the "two camps of scientific psychology," the laboratory scientist and the test constructor, to work together. Almost 40 years later it is clear that there has been a response, and that this work has substantially extended our knowledge of intelligence.

References

Ackerman, P. L. (1988). Determinants of individual differences during skill acquisition: cognitive abilities and information processing. *Journal of Experimental Psychology: General*, **117**, 288–318.

Bell, L. C., & Perfetti, C. A. (1994). Reading skill: some adult comparisons. *Journal of Educational Psychology*, **86**, 244–255.

Carroll, J. B. (1993). *Human cognitive abilities*. Cambridge: Cambridge University Press.

Cattell, R. B. (1971). *Abilities: their structure, growth, and action.* Boston: Houghton Mifflin.

Cerella, J. (1990) Aging and information-processing rate. In J. Birren & K. Schaie (Eds.), *Handbook of the psychology of aging* (3rd ed., pp. 201–222). New York: Academic Press.

Cronbach, L. J. (1957). The two disciplines of scientific psychology. *American Psychologist*, **12**, 671–684.

Daneman, M., & Carpenter, P. A. (1980). Individual differences in working memory and reading. *Journal of Verbal Learning and Verbal Behavior*, **19**, 450–466.

Daneman, M., & Carpenter, P. A. (1983). Individual differences in integrating information between and within sentences. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, **9**, 561–584.

Detterman, D. K., & Daniel, M. H. (1989). Correlations of mental test with each other and with cognitive variables are highest in low IQ groups. *Intelligence*, **13**, 349–360.

di Sessa, A. A. (1993). Toward an epistemology of physics. *Cognition and Instruction*, **10**, 105–225.

Earles, J. A., & Ree, M. J. (1992). The predictive validity of the ASVAB for training. *Educational and Psychological Measurement*, **52**, 721–725.

Gardner, H. (1983). *Frames of mind. The theory of multiple intelligences*. New York: Basic Books.

Guilford, J. P. (1967). *The nature of human intelligence.* New York: McGraw-Hill.

Hartigan, J. A., & Wigdor, A. K. (Eds.) (1989). *Fairness in employment testing: validity, generality, minority issues, and the General Aptitude Test Battery.* Washington: National Academy Press.

Herrnstein, R. J., & Murray, C. (1994). *The bell curve. Intelligence and class structure in American life.* New York: Free Press.

Horn, J. L. (1985). Remodeling old models of intelligence. In B. B. Wolman (Ed.), *Handbook of intelligence. Theories, measurements, and applications* (pp. 267–300). New York: Wiley.

Horn, J. L., & Noll, J. (1994). A system for understanding cognitive capabilities: a theory and the evidence on which it is based. In D. Detterman (Ed.), *Current topics in human intelligence: Vol. 4. Theories of intelligence* (pp. 151–204). Norwood, NJ: Ablex.

Hunt, E. (1995). *Will we be smart enough? A cognitive analysis of the coming workforce.* New York: Russell Sage Foundation.

Hunt, E., & Minstrell, J. (1994). A cognitive approach to the teaching of physics. In K. McGilly (Ed.), *Classroom lessons* (pp. 51–74). Cambridge, MA: MIT Press.

Hunt, E., Frost, N., & Lunneborg, C. E. (1973). Individual differences in cognition: a new approach to intelligence. In G. S. Bower (Ed.), *Advances in learning and motivation* (vol. 7). New York: Academic Press.

Hunt, E., Davidson, J., & Lansman, M. (1981). Individual differences in long term memory access. *Memory and Cognition*, **9**, 599–608.

Hunter, J. E. (1986). Cognitive ability, cognitive aptitudes, job knowledge, and job performance. *Journal of Vocational Behavior*, **29**, 340–362.

Jensen, A. R. (1987). Individual differences in the Hick paradigm. In P. A. Vernon (Ed.), *Speed of information processing and intelligence* (pp. 101–176). Norwood, NJ: Ablex.

Kintsch, W., & Greeno, J. R. (1985). Understanding and solving word arithmetic problems. *Psychological Review*, **92**, 109–129.

Kyllonen, P. C., & Christal, R. E. (1990). Reasoning ability is (little more than) working memory capacity?! *Intelligence*, **14**, 389–433.

Larkin, J. (1983). Problem representation in physics. In D. Gentner & A. Stevens (Eds.), *Mental models.* Hillsdale, NJ: Lawrence Erlbaum Associates.

Lawrence, J. A. (1986). Expertise on the bench: modeling magistrates' judicial decision making. In M. T. H. Chi, R. Glaser, & M. F. Farr (Eds.), *The nature of expertise* (pp. 229–259). Hillsdale, NJ: Lawrence Erlbaum Associates.

McHenry, J. J., Hough, L. M., Toquam, J. L., Hanson, M. A., & Ashworth, S. (1990). Project A validity results: The relationship between predictor and criterion domains. *Personnel Psychology*, **43**, 335–354.

MacLeod, C. M., Dekaban, K. A., & Hunt, E. (1978). Memory impairment in epileptic patients: selective effects of phenobarbital concentration. *Science*, **202**, 1102–1104.

Miller, G. A. (1956). The magical number seven, plus or minus two: some limits on our capacity to process information. *Psychological Review*, **63**, 81–97.

Miyake, A., Just, M. A., & Carpenter, P. A. (1994). Working memory constraints on the resolution of lexical ambiguity: maintaining multiple interpretations in neutral contexts. *Journal of Memory and Language*, **33**, 175–202.

Myerson, J., Hale, S., Wagstaff, D., Poon, L. W., & Smith, G. A. (1990). The information-loss model: a mathematical theory of age-related cognitive slowing. *Psychological Review*, **97**, 475–487.

Nesselroade, J. R., & Thompson, W. W. (1995). Selection and related threats to group comparisons. An example comparing factorial structures of higher and lower ability groups of adult twins. *Psychological Bulletin*, **117**, 271–284.

Palmer, J., MacLeod, C. M., Hunt, E., & Davidson, J. (1985). Information processing correlates of reading. *Journal of Memory and Language*, **24**, 59–88.

Spearman, C. (1927). *The abilities of man.* London: Macmillan.

Sternberg, R. J. (1977). Component processes in analogical reasoning. *Psychological Review*, **84**, 353–378.

Sternberg, R. J. (1990). *Metaphors of mind.* Cambridge: Cambridge University Press.

Sternberg, S. (1966). High speed scanning in human memory. *Science*, **153**, 652–654.

Sternberg, S. (1975). Memory scanning: new findings and current controversies. *Quarterly Journal of Experimental Psychology*, **27**, 1–32.

Thurstone, L. L. (1938). *Primary mental abilities.* Chicago: University of Chicago Press.

Underwood, B. J., Boruch, R. F., & Malmi, R. A. (1978). Composition of episodic memory. *Journal of Experimental Psychology: General*, **107**, 393–419.

Vernon, P. A. (1983). Speed of information processing and general intelligence. *Intelligence*, **7**, 53–70.

Vernon, P. A., & Kantor, L. (1986). Reaction time correlations with intelligence test scores obtained under either timed or untimed conditions. *Intelligence*, **10**, 315–330.

Wigdor, A. K., & Green, B. F., Jr (1991). *Performance assessment in the workplace*. Washington, DC: National Academy Press.

(Received August 23, 1996; accepted August 31, 1996)