

Metacognition, abilities, and developing expertise: What makes an expert student?

ROBERT J. STERNBERG

Department of Psychology, Yale University, Box 208205, New Haven, CT 06520-8205, U.S.A.

Abstract. The main argument of this article is that metacognition is an important part of human abilities, which are, in turn, forms of developing expertise. To the extent that our goal is to understand the bases of individual differences in student academic success, we need to understand metacognition as representing part of the abilities that lead to student expertise, but only as part.

Anyone who may have had the slightest doubt regarding the importance of metacognition to student success need only read the articles in this issue of *Instructional Science* to remove the doubt. The theory and data of Artzt and Armour-Thomas; Carrell, Gajdusek, and Wise; Everson and Tobias; Gourgey; Mayer; and Wolters and Pintrich make a convincing case. The various researchers use a variety of theoretical frameworks, methodologies, subject-matter areas, and arguments to make a fully persuasive case for the importance of metacognition to school success. One could critique any one study or set of results, but the strength of the symposium is in the converging operations, all of which make the identical case.

The importance of metacognition

Because I agree with their fundamental arguments, and believe in the strength of their work, I do not want to fall into the trap of critiquing or arguing about trivia in individual articles: The authors set out to make a case; they made it. Rather, what I would like to do is discuss the role of metacognition in student expertise.

There are a number of truly interesting ideas and findings in the articles in this symposium. I have picked one from each article that I believe especially merits highlighting.

1. Teachers can be loosely classified into three groups, depending on the extent to which they take students' metacognitive functioning into account.

In essence, the groups represent the teachers' own metacognitive functioning. What Artzt and Armour-Thomas refer to as Group X teachers fully take into account students' patterns of metacognitive functioning; Group Y teachers hardly do so at all; and Group Z teachers show a mixed pattern. The especially attractive feature of Table 1 of the Artzt and Armour-Thomas article is that any teacher can read through the table, match his or her own behavior to that of the three groups, and if the teacher wishes, use the table as a basis for conceptual and behavioral change.

As is always true when grouping is involved – whether of students or, as in this case, of teachers, there are dangers of missing fine but potentially important distinctions. With regard to Groups X, Y, and Z, I believe that there are at least three such distinctions that are important.

First, metacognition is diverse. It includes both understanding and control of cognitive processes. Moreover, these constructs are themselves complex. For example, control of cognitive processes includes planning, monitoring, and evaluating activities. A teacher could be high in knowledge of one or several of these aspects, but deficient in others. Assigning an overall group or score might obscure the patterns in the teachers' understandings of students' metacognitive functioning, so that, for example, two teachers who are in Group Y might have very different patterns of understanding.

Second, two issues need to be kept somewhat distinct – understanding of students' metacognition, and knowing how effectively to act upon this understanding. One teacher might understand students' metacognition, but not know how to translate this understanding into effective action. Another teacher might have somewhat less understanding, but a more effective “translation process.”

Third, we need, of course, to remember that metacognition interacts with many other aspects of the student – abilities, personality, learning styles, and so on. A teacher's understanding of metacognition will probably be most useful if it is complemented by an understanding of these other aspects of students' functioning, and of how they interact with metacognition.

2. A generalized conclusion that emerges from the Carrell et al. article is that what matters is not so much what strategies students use in learning to read in a second language, but rather, their knowing when to use these strategies, how to coordinate between strategies, and having a number of different strategies available. In other words, it is metacognition about strategies, rather than the strategies themselves, that appears to be essential.

I strongly endorse this and other conclusions of the Carrell et al. team. But there is one caveat I would wish to add. Much of language learning is automatic, and occurs at the level of implicit learning rather than explicit learning – one is not even aware that the learning is taking place.

When functioning is automatic, metacognitive activity can actually hamper functioning (Sternberg, 1985). For example, many tennis players have had the experience of finding that when they think too much or too deliberately about what they are doing, the quality of their playing declines. An expert typist who starts thinking about where the keys are will type much more slowly. Thus, although metacognitive activities may be quite useful in many aspects of language learning, they are not necessarily always called for. Students need to learn to automatize, which means, in practice, learning to bypass certain conscious metacognitive activity.

3. Everson and Tobias have shown in their research that not only is knowledge important to school success, but so is knowledge about one's knowledge: one's estimation of how much one knows. Those who have experience in teaching know how important this higher-order knowledge is, because when you don't know you don't know something, you are scarcely motivated to learn it!

I would like to add one caveat to the Everson-Tobias findings. Knowledge monitoring always takes place in a context, and with respect to a particular goal or purpose. Consider one's knowledge of a word, say, "repression." The knowledge one would need in order to be considered "knowledgeable" would differ, say, in an introductory-psychology course versus an advanced graduate seminar. The quantity and quality of knowledge one would need for a multiple-choice test would differ from what one would need for an essay test. Thus, monitoring of knowledge must always take into account the context of learning or testing, and the purpose to which the knowledge will be put. The student who successfully monitors what he or she needs to know to get an A on a test in college may or may not be the person who successfully monitors as a professional in a given field. The relationship between the somewhat superficial monitoring that is needed and the deeper monitoring needed by a professional remains to be examined. In general, many successful students become less than successful professionals, and vice versa, so the relationship is probably worth studying at some future time.

4. Gourgey shows that students should not be expected wildly to welcome instruction on metacognitive skills. On the contrary, they may actively resist it, an experience I have had with my own students. When students have become used to and have been rewarded over the years for passive and rather mindless learning, they will not jump at the chance to take a more thoughtful or mindful approach to what they are doing. Often the teacher's greatest challenge is to interest the students in the first place in metacognitive procedures.

The fact that students often do not welcome metacognitive training shows, in my opinion, a failure in our schooling. Students acquire the notion that knowledge is command of a large body of factual data. Metacognitive proce-

dures, in this context, may seem largely irrelevant. Yet what is truly important in life is knowledge for use, not static knowledge that goes no further than demonstration on a rote-recall test. Students need to learn that metacognitive procedures for learning and using information are at least as important as is knowledge of the information to which these procedures are applied.

5. Mayer shows across a variety of problem-solving domains the importance of *metaskill*, of the ability to control and monitor cognitive processes. Part of developing metaskill in students is the students' developing the individual interest, and the teacher's developing the situational interest, that will motivate the students to think about their problem-solving practices.

Schools insufficiently emphasize metaskills, and I believe that one reason is that teachers are themselves uncertain as to just what the metaskills are, or if they are aware of them, of how they should teach them. We need to train teachers explicitly about what the skills are, and how they can be taught.

6. Wolters and Pintrich show that whereas students report differences in levels of motivational and cognitive components of self-regulated learning in different academic contexts, the relations among these components are similar across the various contexts. In other words, the pattern of relation between motivation and cognition is similar across domains. These results suggest that understanding of cognition is neither domain-specific nor domain-general, but domain-specific in some aspects and domain-general in others.

This finding may apply in the domains and with the materials that Wolters and Pintrich studied, but it is probably not as general a finding as it may seem. Generally, students will be motivated to pursue areas where they excel, and will excel in areas where they are motivated, yielding a correlational relation between motivation and cognition. But there are certainly dissociations, as we discover when students are motivated to do some kinds of things in which they do not excel (e.g., the child who wants to be a doctor who can barely get through science courses, or the child who wants to be a lawyer who does not reason well analytically). At the same time, a student may have the ability to do, say, math or science well, but not be interested in it. In these cases, the relation breaks down. The teacher may choose to help try to restore the relation, either by helping the student increase abilities for pursuing a particular field, or by helping a student become more motivated in an area in which he or she excels, but has not been previously motivated.

7. Schraw makes the important distinction between knowledge and regulation of cognition described above, and argues that metacognitive knowledge is multidimensional, domain-general in nature, and teachable. This point of view is optimistic, although perhaps just a bit too optimistic. At one level, it is probably correct. For example, the need to formulate strategies, or to represent information, is domain-general. But the person who is well able

to formulate a strategy or represent information in solving physical-science problems may not be the person who is well able to do the same things in the domain of writing a literary composition, and vice versa.

In our own research (Sternberg, 1985), we have found that content effects can be as large as or larger than process effects on individual differences. Thus, we should not assume that because a given metacognitive or other process has the same name or description across domains, it is equally easy for a given individual to implement across domains, or that patterns of individual differences will be comparable. The best physicists and literary critics, for example, may all be excellent metacognitively, but they might not be in each others' domains, despite the names of the processes (such as strategy formulation) being the same.

What exactly is the role of metacognition, then, in student expertise? And why do we need to discuss the role of metacognition in student expertise? What does it even mean to discuss its "role"? Here's the problem. In the abilities domain, as noted by Spearman (1904) and pretty much everyone since who has studied the problem, everything correlates with everything. The result has been a large, and at times overwhelming, documentation of the role of many things in abilities or in student success (see, e.g., Sternberg (1994), for discussions of many of these factors). Researchers have found, for example, that abilities can be understood in terms of nerve conduction velocity (Vernon & Mori, 1990, 1992), choice reaction time (Jensen, 1993), inspection time (Nettelbeck, 1987), speed of components of inductive reasoning (Sternberg & Gardner, 1983), lexical-decision time (Hunt, 1978, 1987), scores on psychometric factors of abilities (Carroll, 1993), knowledge base and its organization (Chi, Glaser & Rees, 1982), metacognition (Mayer, 1992, and all the articles in this symposium); and many other things. This list only scratches the surface.

The problem is that when there is a positive manifold, almost everything correlates with everything else, and it is easy to slip into causal inferences from these correlations, despite admonitions to the contrary from elementary statistics teachers. Thus, in all of the above cases, arguments have been made not only that these attributes are correlated with abilities or with school success, but that they are somehow causal, necessary, or at least, highly desirable. The last claim is probably the best-supported one. I will argue in this article that metacognition converges with other attributes that have been linked to the abilities necessary for school success in a construct of developing expertise.

School-relevant abilities as forms of developing expertise

The abilities as developing expertise view

How might we best view the abilities necessary for success in school? I would like to argue that these abilities should be viewed in a way that is somewhat distinct from the typical way in which abilities are viewed.

The best available answer to the nature of school-relevant abilities is quite different from the one that is conventionally offered. The argument here is that the scores and the difference between them reflect not some largely inborn, relatively fixed “ability” construct, but rather a construct of developing expertise. I refer to the expertise that all these assessments measure as “developing” rather than as “developed” because expertise is typically not at an endstate, but in a process of continued development.

In a sense, the point of view represented in this article represents only a minor departure from some modern points of view regarding abilities. Abilities are broadly conceived, and are seen as important to various kinds of success. They are seen as modifiable, and as capable of being flexibly implemented. What is perhaps somewhat new here is the attempt to integrate two literatures – the literature on abilities with that on expertise, and to argue that the two literatures may be talking, at some level, about the same thing. Furthermore, metacognition is viewed as part of the concept of developing expertise.

Traditionally, abilities are typically seen either as precursors to expertise (see essays in Chi, Glaser & Farr, 1988) or as opposed to expertise (Fiedler & Link, 1994) as causes of behavior. Sometimes, abilities are held up in contrast to deliberate practice as causes of expertise (Ericsson, Krampe & Tesch-Romer, 1993). Here, abilities are seen as themselves a form of developing expertise. When we test for them, we are as much testing a form of expertise as we are when we test for accomplishments of various kinds, whether academic achievement, skill in playing chess, skill in solving physics problems, or whatever. What differs is the kind of expertise we measure, and more importantly, our conceptualization of what we measure. The difference in conceptualization comes about in part because we happen to view one kind of accomplishment (ability-test scores) as predicting another kind (achievement test scores, grades in school, or other indices of accomplishment). But according to the present view, this conceptualization is one of practical convenience, not of psychological reality.

According to this view, although ability tests may have temporal priority relative to various criteria in their administration, they have no psychological priority. All of the various kinds of assessments are of a kind. What distinguishes ability tests from the other kinds of assessments is how the ability tests

are used, rather than what they measure. There is no qualitative distinction among the various kinds of assessments.

One comes to be an expert in taking ability tests in much the same ways one becomes an expert in anything else – through a combination of genetic endowment and experience.

Characteristics of expertise

The characteristics of experts as reflected in performance on ability tests are similar to the characteristics of experts of any kind (see Chi, Glaser & Farr, 1988; Sternberg, 1996a: 374). For example, operationally, by *expertise*, I refer, in a given domain, to people's: (a) having large, rich schemas containing a great deal of declarative knowledge about a given domain, in the present case, the domains sampled by ability tests; (b) having well-organized, highly interconnected units of knowledge about test content stored in schemas; (c) spending proportionately more time determining how to represent test problems than they do in search for and in executing a problem strategy; (d) developing sophisticated representations of test problems, based on structural similarities among problems; (e) working forward from given information to implement strategies for finding unknowns in the test problems; (f) generally choosing a strategy based on elaborate schemas for problem strategies; (g) having schemas containing a great deal of procedural knowledge about problem strategies relevant in the test-taking domain; (h) having automatized many sequences of steps within problem strategies; (i) showing highly efficient problem solving; when time constraints are imposed, they solve problems more quickly than do novices; (j) accurately predicting the difficulty of solving particular test problems; (k) carefully monitoring their own problem-solving strategies and processes; and (l) showing high accuracy in reaching appropriate solutions to test problems.

Ability tests, achievement tests, school grades, and measures of job performance all reflect overlapping kinds of expertise in these kinds of skills. To do well in school or on the job requires a kind of expertise; but to do well on a test also requires a kind of expertise. Of course, part of this expertise is the kind of test-wiseness that has been studied by Millman, Bishop & Ebel (1965) and others (see Bond & Harman, 1994); but there is much more to test-taking expertise than test-wiseness.

Most importantly from the present point of view, many of the aspects of expertise directly involve metacognition (or what I have called meta-componential functioning [Sternberg, 1985]). For example, in terms of the above list of characteristics of expertise, time allocation (c), development of representations (d), selection of strategies (f), prediction of difficulty (j), and monitoring (k) are all aspects of metacognitive functioning. Thus, metacog-

nition represents an extremely important part of developing expertise, but not the only part, of course. Similarly, aspects of functioning studied by other investigators, such as speed of thinking or organization of knowledge base, also form part of developing expertise in students.

People who are more expert in taking IQ-related tests have a set of skills that is valuable not only in taking these tests, but in other aspects of Western life as well. Taking such a test typically requires metacognitive and other skills such as (a) figuring out what someone else (here, a test constructor) wants, (b) command of English vocabulary, (c) reading comprehension, (d) allocating limited time, (e) sustained concentration, (f) abstract reasoning, (g) quick thinking, (h) symbol manipulation, (i) suppression of anxiety and other emotions that can interfere with test performance, and so on. These skills are also part of what is required for successful performance in school and many kinds of job performance. Thus, an expert test-taker is likely also to have skills that will be involved in other kinds of expertise as well, such as expertise in getting high grades in school.

It is, in my opinion, not correct to argue that the tests measure little or nothing of interest. At the same time, there are many important kinds of expertise that the tests do not measure, for example, what Gardner (1983, 1993) would call musical, bodily-kinesthetic, interpersonal, and intrapersonal intelligences, and what I would call creative and practical intelligences (Sternberg, 1985, 1988).

To the extent that the expertise required for one kind of performance overlaps with the expertise required for another kind of performance, there will be a correlation between performances. The construct measured by the ability tests is not a “cause” of school or job expertise; it is itself an expertise that overlaps with school or job expertise. On the overlapping-expertise view, the traditional notion of test scores as somehow causal is based upon a confounding of correlation with causation. Differences in test scores, academic performance, and job performance are all effects – of differential levels of expertise.

Acquisition of expertise

Individuals gain the expertise to do well on ability tests in much the same way they gain any other kind of expertise – through the interaction of whatever genetic dispositions they bring to bear with experience via the environment. I refer to tests as measuring *developing* expertise because the experiential processes are ongoing. In particular, individuals (a) receive direct instruction in how to solve test-like problems, usually through schooling; (b) engage in actual solving of such problems, usually in academic contexts; (c) engage in role modeling (watching others, such as teachers or other students, solve test-

like problems); (d) think about such problems, sometimes mentally simulating what they might do when confronting such problems; and (e) receive rewards for successful solution of such problems, thereby reinforcing such behavior.

Individual differences in expertise

None of these arguments should be taken to imply that individual differences in underlying capacities do not exist. The problem, as recognized by Vygotsky (1978), as well as many others, is that we do not know how directly to measure these capacities. Measures of the zone of proximal development (e.g., Brown & Ferrara, 1985; Brown & Frensch, 1979; Feuerstein, 1979) seem to assess something other than conventional psychometric *g*, but it has yet to be shown that what it is they do measure is the difference between developing ability and latent capacity.

Individual differences in developing expertise result in much the same way they result in most kinds of learning – from (a) rate of learning (which can be caused by the amount of direct instruction received, amount of problem solving done, amount of time and effort spent in thinking about problems, and so on), and from (b) asymptote of learning (which can be caused by differences in the numbers of schemas, organization of schemas, efficiency in using schemas, and so on). Ultimately, such differences will represent a distinct genetic-environmental interaction for each individual.

There is no evidence, to my knowledge, that individual differences can be wiped out by the kind of “deliberate practice” studied by Ericsson and his colleagues (e.g., Ericsson & Charness, 1994; Ericsson, Krampe & Tesch-Romer, 1993; Ericsson & Smith, 1991). Ericsson’s work shows a correlation between focused practice and expertise; it does not show a causal relation, any more than the traditional work on abilities shows causal relations between measured abilities and expertise. A correlational demonstration is an important one; it is not the same as a causal one.

For example, the fact that experts have tended to show more deliberate practice than novices may itself reflect an ability difference (Sternberg, 1996b). Meeting with success, those with more ability may practice more; meeting with lesser success, those with lesser ability may give up. Or both deliberate practice and ability may themselves be reflective of some other factor, such as parental encouragement, which could lead both to the nurturing of an ability and to practice. Indeed, deliberate practice and expertise may interact bidirectionally, so that deliberate practice leads to expertise, and the satisfaction of expertise leads to more deliberate practice. The point is that a variety of mechanisms might underlie a correlational relationship. It seems unquestionable that deliberate practice plays a role in the development of expertise.

But it also seems extremely likely that its role is as a necessary rather than sufficient condition.

Deliberate practice may play a somewhat lesser role in creative performance than in other kinds of performance (Sternberg, 1996a). We might argue over whether someone who practices memorization techniques can become a mnemonist. Probably, they can become a mnemonist at least within certain content domains (Ericsson, Chase & Faloon, 1980). It seems less plausible that someone who practices composing will become a Mozart. Other factors seem far more important in the development of creative expertise, such as pursuing paths of inquiry that others ignore or dread, taking intellectual risks, persevering in the face of obstacles, and so on (Sternberg & Lubart, 1995, 1996).

Relations among various kinds of expertise

Although all of the various assessments considered here overlap, the overlap is far from complete. Indeed, a major problem with both ability tests and school achievement tests is that the kinds of skills measured in many respects depart from the skills that are needed for job success (see, e.g., Sternberg, Ferrari, Clinkenbeard & Grigorenko, 1996; Sternberg, Wagner, Williams & Horvath, 1995).

An individual can be extremely competent in test and school performance, but flag on the job because of the differences in the kinds of expertise required. For example, success in memorizing a textbook may lead to a top grade in a psychology or education course, but may not predict particularly well whether someone will be an expert researcher or an expert teacher. The creative and practical skills needed for these kinds of job success may be only minimally or not at all tapped in the ability-testing and school-assessment situations. Thus, it is not particularly surprising that although test scores and school grades correlate with job performance, the correlations are far from perfect.

There are various measures that correlate with IQ that do not, on their face, appear to be measures of achievement. But they are measures of forms of developing expertise. For example, the inspection-time task of Nettelbeck (1987; Nettelbeck & Lally, 1976) or the choice reaction-time task of Jensen (1982) both correlate with psychometric *g*. However, performances on both tasks reflect a form of developing expertise, in one case, of perceptual discriminations, in the other case, of quick response to flashing lights or other stimuli.

Temporal priority is not psychological priority

The argument here is that ability tests are typically temporally prior in their measurement to measurements of various kinds of achievements, but what they measure is not psychologically prior. The so-called achievement tests might just as well be used to predict scores on ability tests, and sometimes are, as when school officials attempt to guess college admissions test scores on the basis of school achievement. In viewing the tests of abilities as psychologically prior, we are confounding our own typical temporal ordering of measurement with some kind of psychological ordering. But in fact, our temporal ordering implies no psychological ordering at all. The relabeling of the SAT as the *Scholastic Assessment Test*, rather than *Scholastic Aptitude Test*, reflects this kind of thinking. Nevertheless, the SAT is still widely used as an ability test, and the SAT-II, which measures subject-matter knowledge, as a set of achievement tests.

An examination of the content of tests of intelligence and related abilities reveals that IQ-like tests measure achievement that individuals should have accomplished several years back. Tests such as vocabulary, reading comprehension, verbal analogies, arithmetic problem solving, and the like are all unequivocally tests of achievement. Even abstract-reasoning tests measure achievement in dealing with geometric symbols, skills taught in Western schools. One might as well use academic performance to predict ability-test scores. The problem with regard to the traditional model is not in its statement of a correlation between ability tests and other forms of achievement, but in its proposal of a causal relation whereby the tests reflect a construct that is somehow causal of, rather than merely antecedent to, later success.

Even psychobiological measures (see, e.g., Vernon, 1990) are in no sense “pure” ability measures, because we know that just as biological processes affect cognitive processes, so do cognitive processes affect biological ones. Learning, for example, leads to synaptic growth (Kandel, 1991; Thompson, 1985). Thus, biological changes may themselves reflect developing expertise. In sum, if we viewed tested abilities as forms of what is represented by the term *developing expertise*, then I would have no argument with the use of the term *abilities*. The problem is that this term is usually used in another way – to express a construct that is psychologically prior to other forms of expertise. Such abilities may well exist, but we can assess them only through tests that measure developing forms of expertise expressed in a cultural context. All abilities, including metacognitive ones, are not fixed, but rather forms of developing expertise. Rather than correlating individual aspects of abilities or achievement – such as metacognitive functioning or reaction time or whatever to scores on tests of abilities or achievement – we need to understand all these aspects not as precursors, but as integral elements in the development of

varied forms of expert performance, including those required to achieve high scores on tests.

Acknowledgements

Preparation of this article was supported under the Javits Act program (Grant #R206R50001) as administered by the Office of Educational Research and Improvement, U.S. Department of Education. The opinions expressed in this report do not reflect the positions or policies of the Office of Educational Research and Improvement or the U.S. Department of Education. I am especially grateful to Michel Ferrari, Elena Grigorenko, Joseph Horvath, Todd Lubart, Richard Wagner, and Wendy Williams for collaborations that have led to the development of many of the ideas expressed in this article. Requests for reprints should be sent to Robert J. Sternberg, Department of Psychology, Yale University, Box 208205, New Haven, CT 06520-8205.

References

- Bond, L. & Harman, A.E. (1994). Test-taking strategies. In R.J. Sternberg, ed., *Encyclopedia of Human Intelligence*, Vol. 2 (pp. 1073–1077). New York: Macmillan.
- Brown, A.L. & Ferrara, R.A. (1985). Diagnosing zones of proximal development. In J.V. Wertsch, ed., *Culture, Communication, and Cognition: Vygotskian Perspectives*. New York: Cambridge University Press.
- Brown, A.L. & Frensch, L.A. (1979). The zone of potential development: Implications for intelligence testing in the year 2000. *Intelligence* 3: 255–277.
- Carroll, J.B. (1993). *Human Cognitive Abilities: A Survey of Factor-Analytic Studies*. New York: Cambridge University Press.
- Chi, M.T.H., Glaser, R. & Farr, M., eds. (1988). *The Nature of Expertise*. Hillsdale, NJ: Erlbaum.
- Chi, M.T.H., Glaser, R. & Rees, E. (1982). Expertise in problem solving. In R.J. Sternberg, ed., *Advances in the Psychology of Human Intelligence*, Vol. 1 (pp. 7–75). Hillsdale, NJ: Erlbaum.
- Ericsson, K.A. & Charness, N. (1994). Expert performance: Its structure and acquisition. *American Psychologist* 49: 725–747.
- Ericsson, K.A., Chase, W.G. & Faloon, S. (1980). Acquisition of a memory skill. *Science* 208: 1181–1182.
- Ericsson, K.A., Krampe, R.T. & Tesch-Romer, C. (1993). The role of deliberate practice in the acquisition of expert performance. *Psychological Review* 100: 363–406.
- Ericsson, K.A. & Smith, J., eds. (1991). *Toward a General Theory of Expertise: Prospects and Limits*. New York: Cambridge University Press.
- Feuerstein, R. (1979). *The Learning Potential Assessment Device*. Baltimore, MD: University Park Press.
- Fiedler, F.E. & Link, T.G. (1994). Leader intelligence, interpersonal stress, and task performance. In R.J. Sternberg & R.K. Wagner, eds., *Mind in Context* (pp. 152–167). New York: Cambridge University Press.
- Gardner, H. (1983). *Frames of Mind*. New York: Free Press.
- Gardner, H. (1993). *Multiple Intelligences: The Theory in Practice*. New York: Basic Books.

- Hunt, E.B. (1978). The mechanics of verbal ability. *Psychological Review* 85.
- Hunt, E.B. (1987). The next word on verbal ability. In P.A. Vernon, ed., *Speed of Information Processing and Intelligence* (pp. 347–392). Norwood, NJ: Ablex.
- Jensen, A.R. (1982). The chronometry of intelligence. In R.J. Sternberg, ed., *Advances in the Psychology of Human Intelligence*, Vol. 1 (pp. 255–310). Hillsdale, NJ: Erlbaum.
- Jensen, A.R. (1993). Test validity: *g* versus “tacit knowledge,” *Current Directions in Psychological Science* 1: 9–10.
- Kandel, E. (1991). Cellular mechanisms of learning and the biological basis of individuality. In E.R. Kandel, J.H. Schwartz & T. M. Jessell, eds., *Principles of Neural Science* (3rd ed.). New York: Elsevier.
- Mayer, R.E. (1992). *Thinking, Problem Solving, Cognition* (2nd ed.). New York: W. H. Freeman.
- Millman, J., Bishop, H. & Ebel, R. (1965). An analysis of test-wiseness. *Educational and Psychological Measurement*, 25: 707–726.
- Nettelbeck, T. (1987). Inspection time and intelligence. In P.A. Vernon, ed., *Speed of Information Processing and Intelligence*. Norwood, NJ: Ablex.
- Nettelbeck, T. & Lally, M. (1976). Inspection time and measured intelligence. *British Journal of Psychology* 67: 17–22.
- Spearman, C.E. (1904). ‘General intelligence’ objectively determined and measured. *American Journal of Psychology* 15: 201–293.
- Sternberg, R.J., ed. (1982). *Handbook of Human Intelligence*. New York: Cambridge University Press.
- Sternberg, R.J. (1985). *Beyond IQ: A Triarchic Theory of Human Intelligence*. New York: Cambridge University Press.
- Sternberg, R.J. (1986). *Intelligence Applied*. Orlando, FL: Harcourt Brace College Publishers.
- Sternberg, R.J. (1988). *The Triarchic Mind: A New Theory of Human Intelligence*. New York: Viking-Penguin.
- Sternberg, R.J. (1990). *Metaphors of Mind: Conceptions of the Nature of Intelligence*.
- Sternberg, R.J., ed. (1994). *Encyclopedia of Human Intelligence* (2 vols.). New York: Macmillan.
- Sternberg, R.J. (1995). For whom the bell curve tolls: A review of *The Bell Curve*. *Psychological Science* 6(5): 257–261.
- Sternberg, R.J. (1996a). *Cognitive Psychology*. Orlando: Harcourt Brace College Publishers.
- Sternberg, R.J. (1996b). Costs of expertise. In K.A. Ericsson, ed., *The Road to Excellence* (pp. 347–354). Mahwah, NJ: Lawrence Erlbaum Associates.
- Sternberg, R.J., Conway, B.E., Ketron, J.L. & Bernstein, M. (1981). People’s conceptions of intelligence. *Journal of Personality and Social Psychology* 41: 37–55.
- Sternberg, R.J., Ferrari, M., Clinkenbeard, P. & Grigorenko, E.L. (1996). Identification, instruction, and assessment of gifted children: A construct validation of a triarchic model. *Gifted Child Quarterly* 40: 129–137.
- Sternberg, R.J. & Gardner, M.K. (1983). Unities in inductive reasoning. *Journal of Experimental Psychology: General* 112: 80–116.
- Sternberg, R.J. & Horvath, J.A. (1995). A prototype view of expert teaching. *Educational Researcher* 24(6): 9–17.
- Sternberg, R.J. & Lubart, T.I. (1995). *Defying the Crowd: Cultivating Creativity in a Culture of Conformity*. New York: Free Press.
- Sternberg, R.J. & Lubart, T.I. (1996). Investing in creativity. *American Psychologist* 51: 677–688.
- Sternberg, R.J. & Spear-Swerling, L. (1996). *Teaching for Thinking*. Washington, DC: APA Books.
- Sternberg, R.J., Wagner, R.K., Williams, W.M. & Horvath, J. (1995). Testing common sense. *American Psychologist* 50(11): 912–927.
- Thompson, R.F. (1985). *The Brain: An Introduction to Neuroscience*. New York: Freeman.

- Vernon, P.A. (1990). The use of biological measures to estimate behavioral intelligence. *Educational Psychologist* 25: 293–304.
- Vernon, P.E. (1971). *The Structure of Human Abilities*. London: Methuen.
- Vernon, P.A. & Mori, M. (1990). Physiological approaches to the assessment of intelligence. In C.R. Reynolds & R.W. Kamphaus, eds., *Handbook of Psychological and Educational Assessment of Children: Intelligence and Achievement*. New York: Guilford Press.
- Vernon, P.A. & Mori, M. (1992). Intelligence, reaction times, and peripheral nerve conduction velocity. *Intelligence* 8: 273–288.
- Vygotsky, L.S. (1978). *Mind in Society: The Development of Higher Psychological Processes*. Cambridge, MA: Harvard University Press.