Speech by Richard C. Atkinson on the occasion of being named 2011 Scientist of the Year by the San Diego Chapter of the ARCS Foundation (March 18, 2011).

The Science of Human Memory

When I was asked whether I'd be willing to be named scientist of the year by Achievement Rewards for College Scientists (ARCS), I requested a few days to think about the offer. On the one hand, I have not been actively involved in scientific research for over 30 years, and there are so many brilliant scientists here in San Diego highly deserving of such an honor. On the other hand, it's rare for a behavioral scientist to receive such an award and by accepting I'd be reinforcing the importance of the behavioral sciences in the overall scheme of science. Further, the research that I did in my earlier career has had an enduring impact over these many years, and it would be a pleasant trip down memory lane to describe the work for a San Diego audience. Thus, I accepted the honor with reservations but with great pleasure.

My active research career covered the period from about 1956 to 1975 while I was at Stanford University. I was a faculty member of four departments at Stanford: the psychology department, the department of statistics and applied mathematics, the School of Education, and the department of operations research in the School of Engineering. Most of you will think this is an odd mix of departments, but hopefully you'll see how they are relevant when I describe my research.

Since the middle of the 19th century scientists in Germany and elsewhere in Europe, and later in the United States, have been conducting well-controlled laboratory experiments on the psychology of memory; how individuals memorize information and then retrieve that information from memory at a later time, sometimes seconds or minutes later and sometimes years later. Many of you have probably been subjects in one or another of these experiments. For people in their 40s or older, virtually everyone who took introductory psychology at a major university was required to serve as a subject in several such experiments to receive full credit for the course.

One example of a very simple experiment is what's called *free recall*. The subject first sees a long list of unrelated words presented one at a time. After all the words have been presented, the subject must immediately recall them in any order (hence the designation *free recall*). The result of principal interest is the probability of recalling each word in the list as a function of its place in the list, or "serial presentation position". Plotting this function yields a U-shaped curve as illustrated in the figure below (for a list of 40 words presented at a rate of one per second). The increased level of recall for the first few words presented (on the left of the graph) is called the primary effect: the large increase for the last eight to twelve words (on the right) is called the recency effect. I'll return to this example later.

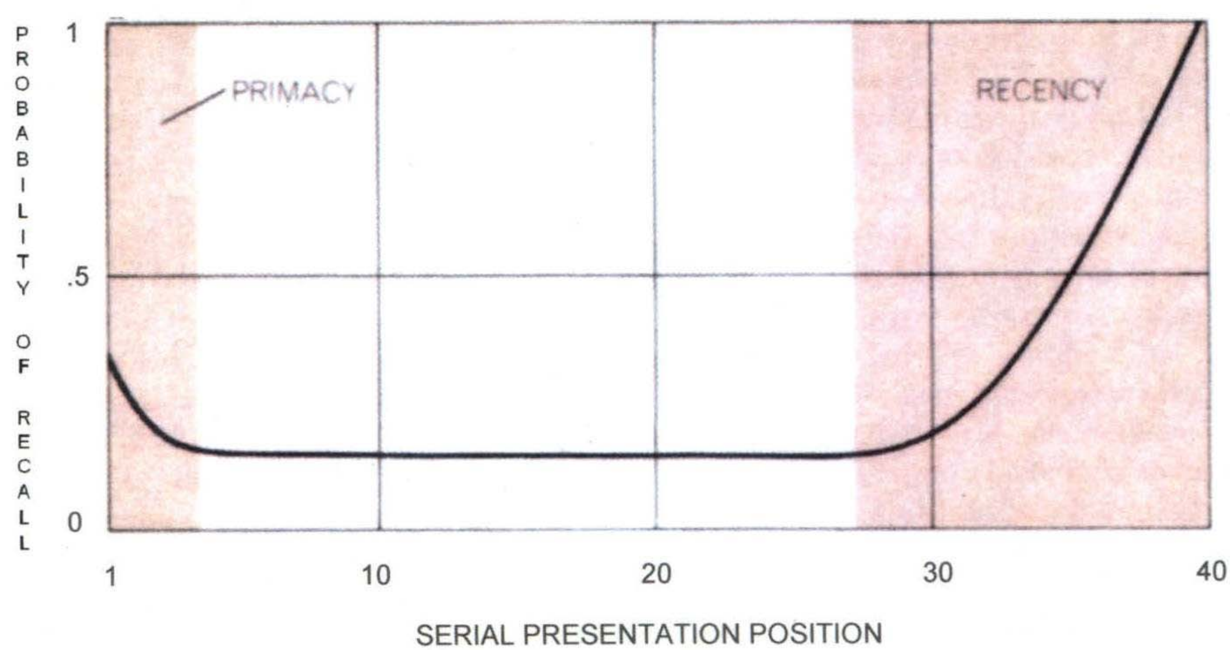

Over the years, an ingenious variety of experiments have been devised probing virtually every aspect of human memory—how information is stored in memory, how information once stored is later retrieved, and the types of activities that can degrade the memory trace or interfere with the retrieval process. The data from these various tasks has proved to be quite orderly and reproducible from one laboratory to another. As the data accumulated for a specific memory task, investigators proposed explanations for what was taking place as the subject memorized and later retrieved the material. These explanations were usually stated in qualitative terms; they provided some insights into the particular phenomenon being studied, but there was little generality from one to the next.

In the early 1960's I began to publish a series of articles laying the framework for a "general theory of human memory", a theory that endeavored to span the full range of experimental findings. The defining publication occurred in 1968 and was titled *Human Memory: A Proposed System and its Control Processes*. My co-author was a young graduate student at Stanford named Richard Shiffrin. I'll have more to say about Rich later.

The theory postulated a system with two components: a short-term memory (STM) and a long-term memory (LTM). The STM is of very limited capacity; its content is continually changing; nothing is stored there on a permanent basis. If you like, think of the contents of STM as a part of our conscious experience—the mental activities of which we are aware. In contrast to STM, LTM is virtually limitless and provides a relatively permanent repository of information, with new information being added over time. Stored in LTM is information about episodes that occur over a lifetime; the knowledge needed to understand and speak a language, and all other information available to us from our memory.

Information from the environment is processed by our various senses (vision, hearing, etc.) and selectively entered into STM. Information once entered into STM can lead to the retrieval of related information in LTM. The retrieved information from LTM helps interpret what is currently in STM, and helps determine if any of the information in STM should be transferred to LTM. A strong assumption of the theory is that the storage of information in LTM must occur via STM. STM is a gatekeeper for what eventually is stored in LTM.

What drives the whole system is a set of control processes that determine what stimuli are attended to, the form in which they are coded in STM and transferred to LTM, the retrieval process for LTM, and the organization of the LTM network. The control processes are too complicated to describe here, but to give you some sense of the system in operation, let me describe a case study of an individual with amnesia of a unique kind.

In 1953, a young man known by his initials as HM (to protect his identity) underwent an experimental operation involving the removal of two finger-shaped slivers of brain tissue, including a structure called the *hippocampus*. The surgeon hoped that the operation would relieve HM's chronic seizures—it proved less than successful. But the operation left HM with profound amnesia of a very special kind. When he was introduced to a new person, he could carry on what appeared to be a normal conversation. However, if he saw the same person again—an hour or a day later—he would have absolutely no recall of having met and talked with the person before. In conversation, he had no trouble recalling events that occurred prior to the operation such as World War II, the Great Depression or the presidency of Franklin Roosevelt. But events that occurred after the operation—the Vietnam War, Nixon's presidency or the demise of the Soviet Union, for example—all came and went without leaving a trace in LTM. For HM, STM continued to work reasonably well as did the retrieval of information from LTM that had been stored there prior to the operation. However, after the operation HM lost all ability to transfer new information to LTM. Over the years, considerable evidence has accumulated to support the idea that the hippocampus plays a critical role in the transfer of information from STM to LTM.

The 1968 Atkinson-Shiffrin paper and a series of related papers laid the foundation for what has come to be known as the *Atkinson-Shiffrin model for memory*. The 1968 paper is one of the most frequently cited articles in the behavioral science literature. Most text books in the field of memory and cognition present the theory and describe it as the "modal model" for a theory of memory. Over the last 40 years, there have been hundreds of articles dealing with one aspect or another of the theory. In 1971, the popular magazine "Scientific American" ran a lead article describing the theory and some of its implications. In 1980, a Russian translation of papers about the theory was published by the USSR Academy of Sciences entitled *Human Memory and the Learning Processes: Selected papers of Richard C. Atkinson*. In 1998 an edited book of articles was published entitled *On Human Memory: Evolution, Progress and Reflections on the 30th Anniversary of the Atkinson-Shiffrin Model*, and in 2003 a special issue of the journal *Cognitive Science* was devoted to honoring the work of Richard Shiffrin and summarizing the current state of the theory. Rich Shiffrin was a young graduate student when we co-authored the 1968 paper. He is now a senior scientist in his sixties who has received a great many honors for his work. I was fortunate to have had the opportunity to work with and mentor an individual who continues to be a highly productive and innovative scientist.

When one examines the history of research on the theory, much of it is based on mathematical models for specific experimental tasks; the models are generally formulated as probabilistic processes and are derived directly from the general theory. The goodness-of-fit of the model's predictions to actual data is a powerful tool for evaluating the theory. The theory has had a considerable impact on the field of memory and cognition because it presented a global framework within which mathematical models could be constructed for many different phenomena. As one textbook author described the theory, it was as if earlier theories "proposed the elements of earth, air, fire and water, and the Atkinson-Shiffrin model proposed the elements found in the periodic table—the latter notion being more complex and comprehensive and explaining a wider variety of phenomena." The theory has changed and evolved over the years, based on new research findings. The term "general theory of human memory" may seem grandiose in describing the 1968 paper, but continuing work from this perspective over these many years has evolved into what I believe is, indeed, a general theory. It is no longer the Atkinson-Shiffrin model, but a theory with many contributors.

As a final remark on the Atkinson-Shiffrin model, let me return to the U-shaped curve described earlier. The mathematical equation for the U-shaped curve is derived directly from the theory, and by estimating three parameters one can use the equation to predict the performance of subjects in the free-recall task. The part of the U-shaped curve to the left in the graph is for the first few words presented and the part to the right is for the last few words. At the time of recall, several of the last few words presented will still be in STM, whereas what traces remain of other words are now in LTM. Recall of the last few words is very high because those words still residing in STM can be readily retrieved. But recall of the first few words is also fairly good. Why? When the first few words are presented, they enter an empty STM; but as more words are presented, STM quickly fills up. At that point, new incoming words will displace older words because of the limited capacity of STM. The first few words reside longer in STM than do later words because they enter an empty STM, and the longer a word resides in STM the more likely a trace of that word will be transferred to LTM. So only the first few items presented enjoyed the extra opportunity of transfer, which is why they were recalled so well.

One of the three parameters that must be estimated to fit the model to data is the capacity of STM; namely, in the free-recall task, the number of unrelated words STM can hold simultaneously. The estimate was approximately 7. In the psychological literature, 7 ± 2 is referred to as the *magical number* because it keeps surfacing in so many different perceptual and cognitive tasks.

The data from several different versions of the free recall task are summarized in the figure on the next page. The top left display is for the basic experiment involving a list of 40 words presented at a one-second rate **(a)**. The lower left display involves a change in the experimental procedure. The subject is required to carry out a difficult arithmetic task for 30 seconds immediately following presentation of the list and then is asked to recall. One can assume that the arithmetic task causes the loss of all the words in STM so that recall reflects retrieval from LTM only. The recency effect is eliminated when this experiment is performed; the earlier portions of the serial-position curve are unaffected **(b)**. If variables that influence the long-term store but not the short-term store are manipulated, the recency portion of the serial-position curve should be relatively unaffected, whereas the earlier portions of the curve should show changes. One such variable is the number of words in the presented list. A word in a longer list is less likely to be recalled, but the recency effect is quite unaffected by list length **(c)**. Similarly, increases in the rate of presentation decrease the likelihood of recalling words preceding the recency region but leave the recency effect largely unchanged **(d)**.

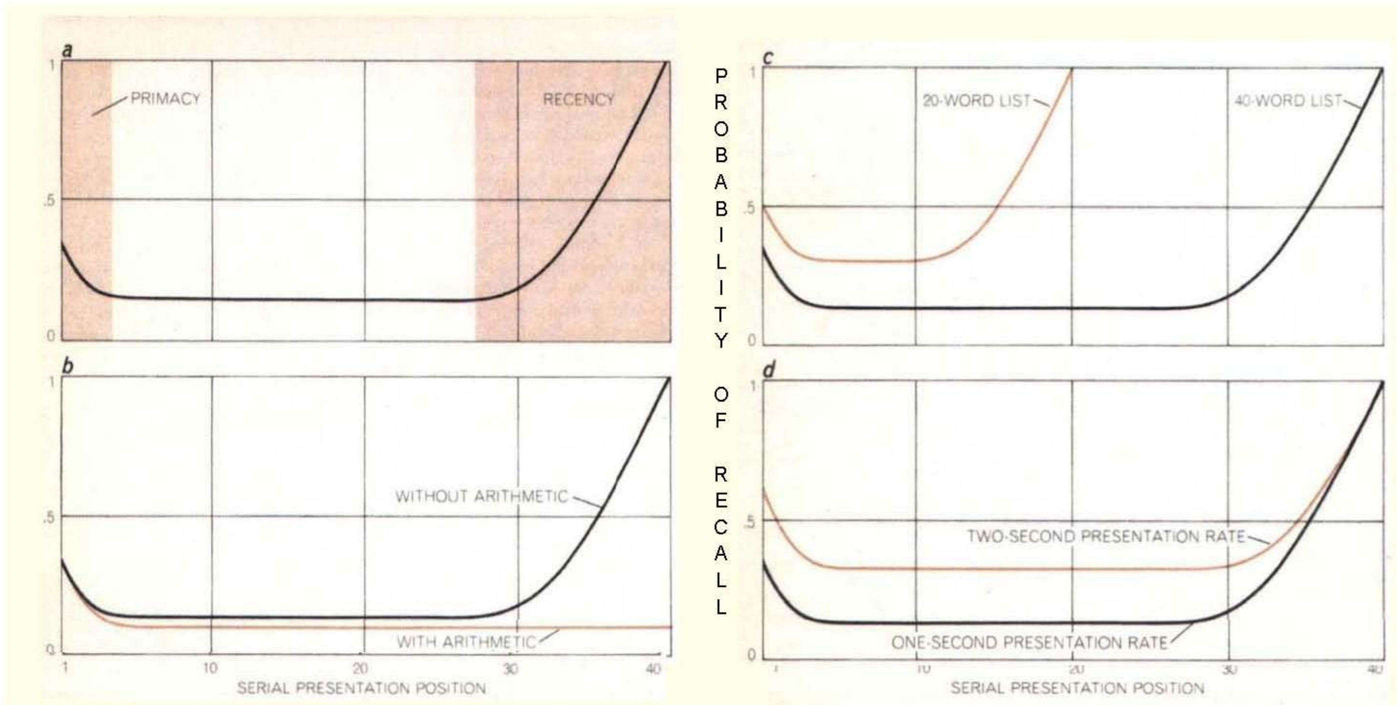

PROBABILITY OF RECALL in free-recall experiments varies in a characteristic way with an item's serial position in a list: a "primary effect" and a "recency effect" are apparent (a). If an arithmetic task is interpolated between presentation and recall, the recency effect disappears (b). Words in long lists are recalled less well than words in short lists (c). Slower presentation also results in better recall (d).

The data from these various experiments are predicted with great accuracy by the model, using the same three parameter estimates for all four of the experimental conditions. Predictions about other aspects of the data, such as the order of recall, can be made but that is beyond the scope of my remarks tonight. Hopefully, the figure illustrates the generality of the model across experimental conditions.

The free-recall task is an example of a very simple memory experiment. The model is readily stated in mathematical terms and predictions can be derived using formal mathematical methods. However, that is not always the case, particularly for experiments involving very complicated memory tasks. Even for these experiments one can usually formulate the model mathematically, but the model may be so complex that predictions cannot be derived using standard mathematical methods. In these cases, the model can be restated in computational terms and predictions generated by computer simulation. With modern computers, such simulations are feasible, even for the most complicated models.

You'll recall that I said I was a member of four departments at Stanford. Why four? The psychology department was the focal point for the study of human memory. The theoretical work was heavily mathematical and required statistical expertise, hence the department of statistics and applied mathematics. While at Stanford, I became interested in applying my ideas on memory and cognition to the practical problem of teaching young children in grades K-3. A colleague and I developed a *computer-assisted instruction system* for teaching reading and mathematics to youngsters in East Palo Alto and New York City. The sequence of instruction was controlled by the computer and was highly individualized to take account of each student's history of correct and incorrect responses. Our computer system was the first of its kind and attracted international attention. The work on computer-assisted instruction involved graduate students and faculty from both the School of Engineering and the School of Education. During my years at Stanford, I had Ph.D. students from all four departments.

In 1975, I left Stanford to go to the National Science Foundation (NSF). At the time, I thought I would be in Washington, DC for two years and would then return to Stanford to reengage in teaching and research. It didn't work out that way. How I was able to walk away from my earlier career so easily mystifies me even to this day.

Once at NSF, my attention turned to issues of science policy. There is one accomplishment during that period that I'd like to describe as a conclusion to my remarks tonight. In the 1970s there was great concern about America's competitiveness. A reinvigorated Europe and Japan, declining growth in U.S. productivity, and rising unemployment made economic competitiveness a major national preoccupation. Research universities were producing a rich array of potentially useful research, but innovations were not moving from university laboratories into the private sector as quickly and efficiently as the economy required.

As director of NSF, I organized a group of colleagues to embark on a series of policy studies designed to rebuild the nation's competitiveness. Our proposals included diverse actions such as tax credits for corporate investments in research & development, funding joint university/industry research centers, easing antitrust regulations to encourage more cooperative research between universities and corporations, and giving federal grants to small businesses involved in innovative research. The administration and congress enacted several pieces of legislation that grew out of our policy work. The proposal that I am particularly proud of is the Bayh-Dole Act of 1980. Prior to that legislation, the federal government owned the rights to any patentable discovery coming out of university research funded by NSF, NIH, DARPA and other federal agencies. Yet few research results ever made it from the university laboratory to the marketplace under this arrangement. The Bayh-Dole Act transferred the government's patent rights to universities, leaving it to each university to decide how to share income derived from a patent between the individual researcher and the university. I helped recruit Senator Evan Bayh (D-Indiana) and Senator Robert Dole (R-Kansas) to carry the legislation, and the NSF general counsel prepared a draft of the legislation. Although the result was to open a new income stream to universities, this was secondary to Bayh-Dole's primary aim: namely, to see that the public investment in basic research served the national economy.

The Bayh-Dole Act was critical in defining the research university system that the United States has in place today, with its key role in driving new industries and technologies. As chancellor of UCSD and president of the University of California, two of my goals were to ensure the quality of university research, but also to ensure the effective transfer of the research into the private sector. The Bayh-Dole Act made those goals attainable.

Let me end with a few of the lessons I have learned from my experiences:

- Graduate students are critically important. I have given you one example from my own research. What Rich Shiffrin and I were able to accomplish together is constantly replicated in universities today. Simply put, graduate students are indispensable to the scientific enterprise. The kind of graduate education we have developed in the United States—in which graduate students learn from faculty while having the chance to make valuable contributions of their own—is unsurpassed anywhere in the world.
- Organizations like ARCS make a tremendous contribution to the future of science by supporting, encouraging, and recognizing scientific talent. Anyone who cares about the future of science is in your debt.

- Science is a wonderful life's work and never runs out of questions that need answers. The students here tonight will not lack for fascinating questions to pursue. As one of my scientific heroes Vannevar Bush—science advisor to President Roosevelt during World War II—once wrote, science is an endless frontier.