

Mind as Computer: Birth of a Metaphor

Gerd Gigerenzer and **Daniel G. Goldstein**

Center for Adaptive Behavior and Cognition,
Max Planck Institute for Psychological Research, Munich
and
Department of Psychology,
University of Chicago

Abstract. Scientific discovery has long been explained in terms of theory, data, and little else. We propose a new approach to scientific discovery in which tools play a central role by suggesting themselves as scientific theories, by way of what we call the tools-to-theories heuristic of scientific discovery. In this article, we extend our previous analysis of statistical tools that became theories of mind to the computer and its impact on psychological theorizing. We first show how a conceptual separation of intelligence and calculation in the early 19th century made mechanical computation, and later the electronic computer, conceivable. We next show how in this century, when computers finally became standard laboratory tools, the computer was proposed—and eventually adopted—as a model of mind. Thus, we travel the full circle from mind to computer and back.

I believe that at the end of the century the use of words and general educated opinion will have altered so much that one will be able to speak of machines thinking without expecting to be contradicted. (Alan Turing, 1950, p. 442)

What role do tools play in scientific discovery? Much of 20th-century Anglo-American philosophy of science has revolved around chicken-and-egg debates concerning theory and data and the relation between the two. Until very recently, few philosophers of science have looked at scientific tools, scientific practice, and what actually happens in the laboratory. Tools played no role for the logical empiricist mainstream, for its contemporary variants, or even for those who have opposed mainstream thought, such as Paul Feyerabend and N. R. Hanson. Have philosophers of science spent too little time inside the laboratories to be drawn in by the glamour of technology? Tools, after all, fascinate scientists. Only recently, scholars have moved from looking at science through the lens of theory and data toward an integrated view that considers daily laboratory practice (e.g., Danziger, 1990; Galison, 1987; Gooding, Pinch, & Schaffer, 1989; Hacking, 1983).

We propose a new approach to scientific discovery in which tools play a central role. The classical view of tools in scientific discovery is that they generate new data, which in turn can lead to new theories. But our thesis depicts tools as having an even more active role in discovery. We argue that new tools can directly, rather than through new data, inspire new theories. Consider when the mechanical clock became known to the public, and later the indispensable tool for astronomical research. The universe itself became understood as a kind of mechanical clock with God as the divine watchmaker. Tools become theories.

In this article, we argue that scientists tend to make analogies between new tools and nature, society, or mind. We focus on analytical tools such as statistical techniques and computer programs rather than on mechanical tools such as clocks or microscopes.

Tools-To-Theories Heuristic

We speak of scientific discovery in terms of heuristics of discovery. Our thesis of a *tools-to-theories heuristic* is twofold (Gigerenzer, 1991, 1992, 1994; Gigerenzer & Murray, 1987):

- (1) *Discovery*. New scientific tools, once entrenched in a scientist's daily practice, suggest new theoretical metaphors and concepts.
- (2) *Acceptance*. Once proposed by an individual scientist (or a group), the new theoretical metaphors and concepts are more likely to be accepted by the scientific community if the members of the community are also users of the new tools.

The discoveries and new ideas we deal with in this article are not those of a Newton or Darwin, but they rank among the best that cognitive psychology has produced. By way of the tools-to-theories heuristic, these theories have changed the conception of the mind with their analogies to tools.

Gigerenzer (1991, 1992, 1994; Gigerenzer & Murray, 1987) documented how tools of statistical inference, once institutionalized in experimental psychology as the *sine qua non* of scientific method in the 1950s, turned into theories of mind. For instance, one of the most widely used tools for statistical inference is analysis of variance (ANOVA). By the late 1960s, about 70% of all experimental articles in psychological journals already used ANOVA (Edgington, 1974). The tool became a theory of mind. In his causal attribution theory, Kelley (1967) postulated that the mind attributes a cause to an effect in the same way as behavioral scientists do—namely, by performing an ANOVA. Psychologists were quick to accept the new analogy between mind and their laboratory tool. In the 1970s, attribution theory defined mainstream social psychology; Kelley and Michaela (1980) counted more than 900 references to their theory in one decade.

The new conception of the mind as an intuitive statistician radically changed psychological theorizing. Causal perception became modeled after the tool—more precisely, after the practical use of the tool in psychological laboratories. One practical use of the tool is making inductive inferences from data to hypotheses. Consequently, causal perception became seen as data-driven inductive inference rather than, as in the early work of Michotte (1946/1963), a direct and spontaneous process that requires no inference. Similarly, because ANOVA is only used for one kind of causal inference, the meaning of causal attribution was reduced from the 4 Aristotelian causes in Michotte's work and from the 17 causes that Piaget (1930) distinguished in children's minds to the 1 cause for which ANOVA is used as a tool. Kelley's (1967) ANOVA theory is but one instance of a tools-to-theories discovery in psychology. Conceptual transformation by the analogy with the statistical tool happened independently in several other reaches of psychology, from sensory discrimination to shape perception to recognition memory, among others (Gigerenzer & Murray, 1987).¹

¹ Although we are only dealing with theories of mind, this does not imply that the tools-to-theories heuristic is not applicable in the analysis of other scientific domains (Gigerenzer et al., 1989). Schaffer (1992) provided several examples from the history of electromagnetism, in which theories stemmed from tools. For instance, in 1600, the court physician William Gilbert described the Earth as a vast spherical magnet. This new idea stemmed from the tool he had invented (a magnet, the small *terrella*) and subsequently used as an analogy to understand the world. This projection had consequences. Gilbert inferred that, because his *terrella* rotated, so did the Earth. The tool proved Copernicanism. Lenoir (1986) traced how Faraday's instruments for recording electric currents shaped the understanding of electrophysiological processes by promoting concepts such as "muscle current" and "nerve current."

In this article, we extend our previous analysis from techniques of statistical inference to another tool, the computer. This article is divided into two parts. In the first part, we argue that a conceptual divorce between intelligence and calculation circa 1800, motivated by economic transformations, made mechanical computation (and ultimately the computer) conceivable. The tools-to-theories heuristic comes into play in the second part, in which we show how, when computers finally became standard laboratory tools in this century, the computer was proposed, and with some delay accepted, as a model of mind. Thus, we travel in a full circle from mind to computer and back.

From Mind to Computer

Well, Babbage, what are you dreaming about?” to which I replied, “I am thinking that all these tables (pointing to the logarithms) might be calculated by machinery. (Charles Babbage, c. 1812/1994, p. 31)

The president of the Astronomical Society of London, Henry Colebrooke (1825), summed up the significance of Babbage’s work: “Mr. Babbage’s invention puts an engine in place of the computer” (p. 510). This seems a strange statement about the man who is now praised for having invented the computer. But, at Babbage’s time, the computer was a human being—in this case, someone who was hired for exhaustive calculations of astronomical and navigational tables.

How did Babbage (1791–1871) ever arrive at the idea of putting a mechanical computer in place of a human one? A divorce between intelligence and calculation, as Daston (1994) argued, made it possible for Babbage to conceive this idea.

In the Enlightenment, calculation was not considered a rote, mechanical thought process. In contrast, philosophers of the time held that intelligence and even moral sentiment were in their essence forms of calculation (Daston, 1988, 1994). Calculation was the opposite of the habitual and the mechanical, remote from the realm of menial labor. For Condillac, d’Alembert, Condorcet, and other Enlightenment philosophers, the healthy mind worked by constantly taking apart ideas and sensations into their minimal elements, then comparing and rearranging these elements into novel combinations and permutations. Thought was a combinatorial calculus, and great thinkers were proficient calculators. In the eulogies of great mathematicians, for instance, prodigious mental reckoning was a favorite topic—Gauss’ brilliant arithmetic was perhaps the last of these stock legends. Calculation was the essence of moral sentiment, too. Even self-interest and greed (as opposed to dangerous passions), by their nature of being calculations, were at least predictable and thereby thought to reinforce the orderliness of society (Daston, 1994).

The Computer as a Factory of Workers

By the turn of the 19th century, calculation was shifting from the company of *hommes éclairés* and savants to that of the unskilled work force. Extraordinary mental arithmetic became associated with the idiot savant and the sideshow attraction. Calculation became seen as dull, repetitive work, best performed by patient minds that lacked imagination. Women ultimately staffed the “bureaux de calculs” in major astronomical and statistical projects (despite their earlier being accused of vivid imaginations and mental restlessness; see Daston, 1992). Talent and genius ceased to be virtuoso combinatorics and permutations and turned into romantic, unanalyzable creations. Thereby, the stage became set for the neo-romanticism in 20th-century philosophy of science that declared

creativity as mystical and the context of discovery as “irrelevant to the logical analysis of scientific knowledge” (Popper, 1959, p. 31).

Daston (1994) and Schaffer (1994) argued that one force in this transformation was the introduction of large-scale division of labor in manufacturing, as evidenced in the automatic system of the English machine-tool industry and in the French government’s large-scale manufacturing of logarithmic and trigonometric tables for the new decimal system in the 1790s. French engineer Gaspard Riche de Prony organized the French government’s titanic project for the calculation of 10,000 sine values to the unprecedented precision of 25 decimal places and some 200,000 logarithms to 14 or 15 decimal places during the French Revolution. Inspired by Adam Smith’s praise of the division of labor, Prony organized the project in a hierarchy of tasks. At the top was a handful of excellent mathematicians, including Adrien Legendre and Lazare Carnot, who devised the formulae; in the middle were 7 or 8 persons trained in analysis; at the bottom were 70 or 80 unskilled persons who knew only the rudiments of arithmetic and who performed millions of additions and subtractions. These “manufacturing” methods, as Prony called them, pushed calculation away from intelligence and toward work. The terms *work* and *mechanical* were linked both in England and in France until the middle of the 19th century (Daston, 1994). Work concerned the body but not the mind; in large-scale manufacturing, each worker did only one thing his or her whole life.

After it was shown that elaborate calculation could be carried out by an assemblage of unskilled workers, each knowing very little about the large computation, it became possible for Babbage to conceive of replacing these workers with machinery. Babbage’s view of the computer bore a great resemblance to a factory of unskilled human workers. When Babbage talked about the parts of his “Analytical Engine,” the arithmetic computation and the storage of numbers, he called these the “mill” and the “store,” respectively (Babbage, 1994, p. 23). The metaphor came from the textile industry, in which yarns were brought from the store to the mill, were woven into fabric, and were then sent back to the store. In the Analytical Engine, numbers were brought from the store to the arithmetic mill for processing, and the results were returned to the store. Commenting on this resemblance, Lady Lovelace said, “We may say most aptly that the Analytical Engine weaves algebraic patterns just as the Jacquard loom weaves flowers and leaves” (Babbage, 1994, p. 27).² In his chapter on the “division of mental labor,” Babbage explicitly referred to the French government’s program for the computation of new decimal tables as the inspiration and foundation of a general science of machine intelligence.

Let us summarize the argument. During the Enlightenment, calculation was the distinctive activity of the scientist and the genius and the very essence of the mental life. New ideas and insights were assumed to be the product of the novel combinations and permutations of ideas and sensations. In the first decades of the 19th century, numerical calculation was separated from the rest of intelligence and demoted to one of the lowest operations of the human mind. After calculation became the repetitive task of an army of unskilled workers, Babbage could envision mechanical computers replacing human computers. Pools of human computers and Babbage’s mechanical computer manufactured numbers in the same way as the factories of the day manufactured their goods.³

² The Jacquard loom, a general-purpose device loaded with a set of punched cards, could be used to weave infinite varieties of patterns. Factories in England were equipped with hundreds of these machines, and Babbage was one of the “factory tourists” of the 1830s and 1840s.

³ Calculation became dissociated and opposed not only to the human intellect but also to moral impulse. Madame de Staël, for instance, used the term *calcul* only in connection with the “egoism and vanity” of those opportunists who exploited the French Revolution for their own advantage and selfishness (Daston, 1994).

The Computer as a Brain

Babbage's dream that all tables of logarithms could be calculated by a machine, however, did not turn into a reality during his lifetime. He never completed any of the three machines he had started to build. Modern computers, such as the ENIAC and the EDVAC at the University of Pennsylvania, came about during and after World War II. Did the fathers of computer science see the mind as a computer? We argue that the contemporary analogy stating that the mind is a computer was not yet established before the "cognitive revolution" of the 1960s. As far as we can tell, two groups were willing to draw a parallel between the human and the computer, but neither used the computer as a theory of mind. One group, which tentatively compared the nervous system and the computer, is represented by Hungarian mathematician John von Neumann (1903–1957). The other group, which investigated the idea that machines might be capable of thought, is represented by English mathematician and logician Alan Turing (1912–1954).

Von Neumann (1958), known as the father of the modern computer, wrote about the possibility of an analogy between the computer and the human nervous system. It seems that von Neumann's reading of Warren McCulloch and Walter Pitts' (1943) paper, "A Logical Calculus of the Ideas Immanent in Nervous Activity" triggered his interest in information processing in the human brain soon after the paper was published (Aspray, 1990). McCulloch and Pitts' paper starts with the statement that, because of the all-or-none character of the nervous system, neural events can be represented by means of propositional logic. The McCulloch-Pitts model did not deal with the structure of neurons, which were treated as "black boxes." The model was largely concerned with the mathematical rules governing the input and output of signals. In a 1945 report on EDVAC (the Electronic Discrete Variable Computer), von Neumann described the computer as being built from McCulloch and Pitts' idealized neurons rather than from vacuum tubes, electromechanical relays, or mechanical switches. Understanding the computer in terms of the human nervous system appeared strange to many, including the chief engineers of the ENIAC project, Eckert and Mauchly (Aspray, 1990, p. 173). But, von Neumann hoped that his theory of natural and artificial automata would improve understanding of the design both of computers and of the human nervous system. His last work (for the Silliman lectures), neither finished nor delivered due to illness, was largely concerned with pointing out similarities between the nervous system and the computer—between the neuron and the vacuum tube—but added cautionary notes on their differences (von Neumann, 1958).

What was the reception of von Neumann's tentative analogy between the nervous system and the computer? His intellectual biographer, Aspray (1990, p. 181), concluded that psychologists and physiologists were less than enthusiastic about the McCulloch-Pitts model; Seymour Papert spoke of "a hostile or indifferent world" (McCulloch, 1965, p. xvii), and McCulloch himself admitted the initial lack of interest in their work (p. 9).

The Computer as a Mind

Von Neumann and others looked for a parallel between the machine and the human on the level of hardware. Turing (1950), in contrast, thought the observation that both the modern digital computer and the human nervous system are electrical was based on a "very superficial similarity" (p. 439). He pointed out that the first digital computer, Babbage's Analytical Engine, was purely mechanical (as opposed to electrical) and that the important similarities to the mind are in function rather than in hardware. Turing discussed the question of whether machines can think rather than the question of whether the mind is like a computer. Thus, he was looking in a direction *opposite*

that in which psychologists were looking after the cognitive revolution, and consequently he did not propose any theories of mind. For example, the famous Turing test is about whether a machine can imitate a human mind but not vice versa. Turing argued that it would be impossible for a human to imitate a computer, as evidenced by the human's inability to perform complex numerical calculations quickly. Turing also discussed the question of whether a computer could be said to have free will, a property of humans. Many years later, cognitive psychologists, under the assumptions that the mind is a computer and that computers lack free will, pondered the question of whether humans could be said to have free will. A similar story to this is that Turing (1947/1969) contemplated teaching machines to be intelligent using the same principles used to teach children. The analogy of the computer as a mind was reversed again after the cognitive revolution, as McCorduck (1979) pointed out, when Massachusetts Institute of Technology (MIT) psychologists tried to teach children with the very methods that had worked for computers.

Turing (1947/1969) anticipated much of the new conceptual language and even the very problems Allen Newell and Herbert Simon later attempted to address, as we see in the second part of this article. With amazing prophecy, Turing suggested that nearly all intellectual issues can be translated into the form "Find a number n such that ..."; that is, he suggested that searching is the key concept for problem solving and that Whitehead and Russell's (1935) *Principia Mathematica* might be a good start for demonstrating the power of the machine (McCorduck, 1979, p. 57).

Not only did Turing's life end early and under tragic circumstances, but his work had practically no influence on artificial intelligence in Britain until the mid-1960s (McCorduck, 1979, p. 68). Neither von Neumann nor his friends were persuaded to look beyond similarities between cells and diodes to functional similarities between humans and computers.

To summarize, we have looked at two groups who compared humans and computers before the cognitive revolution. One of these groups, represented by von Neumann, spoke tentatively about the computer as a brain but warned about taking the analogy too far. The other group, represented by Turing, asked whether the computer has features of the human mind but not vice versa—that is, this group did not attempt to design theories of mind through the analogy of the tool.

Before the second half of the century, the mind was not yet a computer. However, a new incarnation of the Enlightenment view of intelligence as a combinatorial calculus was on the horizon.

From Computer to Mind

The computer is a member of an important family of artifacts called symbol systems, or more explicitly, physical symbol systems. ... The hypothesis is that a physical symbol system ... has the necessary and sufficient means for general intelligent action. (Herbert Simon, 1969, p. 26)

What has been called in retrospect the *cognitive revolution* in American psychology of the 1960s is more than an overthrow of behaviorism by mentalist concepts. The cognitive revolution did more than revive the mental; it changed its meaning. One source of this change is the projection of new tools (i.e., statistics and computers) into the mind. We refer to this heuristic of discovery as the *tools-to-theories heuristic*. The two new classes of theories that emerged and that partially overlap pictured the new mind as an "intuitive statistician" or a "computer program."

In this section, we see how a tools-to-theories explanation accounts for the new conception of the mind as a computer, focusing on the discovery and acceptance of Simon and Newell's brand of information-processing psychology. We try to reconstruct the discovery of Newell and Simon's (1972) information-processing model of mind and its (delayed) acceptance by the psychological community in terms of the tools-to-theories heuristic.

Discovery

Babbage's mechanical computer was preceded by human computers. Similarly, Newell and Simon's first computer program, the "Logic Theorist" (LT) was preceded by a human computer. Before the LT was up and running, Newell and Simon reconstructed their computer program out of human components (*viz.*, Simon's wife, children, and several graduate students) in order to see if it would work. Newell wrote up the subroutines of the LT program on index cards:

To each member of the group, we gave one of the cards, so that each person became, in effect, a component of the LT computer program—a subroutine—that performed some special function, or a component of its memory. It was the task of each participant to execute his or her subroutine, or to provide the contents of his or her memory, whenever called by the routine at the next level above that was then in control. So we were able to simulate the behavior of the LT with a computer consisting of human components. ... The actors were no more responsible than the slave boy in Plato's *Meno*, but they were successful in proving the theorems given them. (Simon, 1991, p. 207)

The parallels to Prony's *bureaux de calculs* and the large-scale manufacturing of the new factories of the early 19th century are striking. At essence is a division of labor, in which the work is done by a hierarchy of humans—each requiring little skill and repeating the same routine again and again. Complex processes are achieved by an army of workers who never see but a little piece of the larger picture.⁴

However, between Prony's human computer and Simon's human computer is an important difference. Prony's human computer and Babbage's mechanical computer (modeled after it) performed numerical calculations. Simon's human computer did not. Simon's humans matched symbols, applied rules to symbols, and searched through lists of symbols—in short, performed what is now generally known as symbol manipulation.

The reader will recall from the first part of this article that the divorce between intelligence and numerical calculation made it possible for Babbage to replace the human computer with a mechanical one. In the 20th century, intelligence and calculation are still divorced. Given this divorce and the early conception of the computer as a fancy number cruncher, it is no wonder that the computer never suggested itself as a theory of mind. We argue that an important precondition for the view of mind as a computer is the realization that computers are symbol-manipulation devices in addition to being numerical calculators. Newell and Simon were among the first to realize this. In interviews with Pamela McCorduck (1979), Newell recalled, "I've never used a computer to do any numerical processing in my life" (p. 129). Newell's first use of the computer at RAND Corporation—a prehistoric card-programmed calculator hooked up to a line printer—was printing symbols representing airplanes for each sweep of a radar antenna.

The symbol-manipulating nature of the computer was important to Simon because it corresponded to some of his earlier views on the nature of intelligence:

The metaphor I'd been using, of a mind as something that took some premises and ground them up and processed them into conclusions, began to transform itself into a notion that a mind was something which

⁴ The Manhattan Project at Los Alamos, where the atomic bomb was constructed, housed another human computer. Although the project could draw on the best technology available, in the early 1940s mechanical calculators (e.g., the typewriter-sized Marchant calculator) could only add, subtract, multiply, and, with some difficulty, divide. Richard Feynman and Nicholas Metropolis arranged a pool of people (mostly scientists' wives, who were getting paid three eighths of the scientists' salary), each of whom repetitively performed a small calculation (e.g., cubing a number) and passed the result on to another person, who incorporated it into yet another computation (Gleick, 1992).

took some program inputs and data and had some processes which operated on the data and produced output. (cited in McCorduck, 1979, p. 127)

It is interesting to note that 20 years after seeing the computer as a symbol-manipulating device, Newell and Simon came forth with the explicit hypothesis that a physical symbol system is necessary and sufficient for intelligence.

The Logic Theorist generated proofs for theorems in symbolic logic—specifically, the first 25 or so theorems in Whitehead and Russell’s (1935) *Principia Mathematica*. It even managed to find a proof more elegant than the corresponding one in the *Principia Mathematica*.

In the summer of 1958, psychology was given a double-dose of the new school of information-processing psychology. One dose was the publication of Newell, Shaw, and Simon’s (1958) *Psychological Review* article, “Elements of a Theory of Human Problem Solving”; the other dose was the Research Training Institute on the Simulation of Cognitive Processes at the RAND Corporation, which we discuss later.

The *Psychological Review* article is an interesting document of the transition between the view that the LT is a tool for proving theorems in logic (the artificial intelligence view) and an emerging view that the LT is a model of human reasoning (the information-processing view). In fact, Newell et al. (1958) went back and forth between both views, explaining that “the program of LT was not fashioned directly as a theory of human behavior; it was constructed in order to get a program that would prove theorems in logic” (p. 154); later, they wrote that the LT “provides an explanation for the processes used by humans to solve problems in symbolic logic” (p. 163). The evidence provided for projecting the machine into the mind is mainly rhetorical. For instance, Newell et al. spent several pages arguing for the resemblance between the methods of LT and concepts (e.g., “set,” “insight,” “hierarchy”) described in the earlier psychological literature on human problem solving.

In all fairness, despite Newell et al.’s claim, the resemblance to these earlier concepts as they were used in the work of Karl Duncker, Wolfgang Köhler, and others is slight. New discoveries, by definition, clash with what has come before, but it is often a useful strategy to hide the amount of novelty and to claim historical continuity. When Tanner and Swets (1954) proposed (in the *Psychological Review* four years earlier) that another scientific tool (i.e., the Neyman-Pearsonian techniques of hypothesis testing) would model the cognitive processes of stimulus detection and discrimination, their signal-detection model also clashed with earlier notions, such as the notion of a sensory threshold. Tanner and Swets, however, chose not to conceal this schism between the old and the new theories, explicitly stating that their new theory “appears to be inconsistent with the large quantity of existing data on this subject” (p. 401). As we argued before, there is a different historical continuity in which Newell and Simon’s ideas stand—the earlier Enlightenment view of intelligence as a combinatorial calculus.

Conceptual Change

Newell et al. (1958) tried to emphasize the historical continuity of what was to become their new information-processing model of problem solving, as did Miller, Galanter, and Pribram (1960) in their *Plans and the Structure of Behavior* when they linked their version of Newell and Simon’s theory to many great names such as William James, Frederic Bartlett, and Edward Tolman. We believe that these early claims for historical continuity served as protection: George Miller, who was accused by Newell and Simon as having stolen their ideas and gotten them all wrong, said, “I had to put the scholarship into the book, so they would no longer claim that those were their ideas.

As far as I was concerned they were old familiar ideas” (Baars, 1986, p. 213). In contrast to this rhetoric, here we emphasize the discontinuity introduced by the transformation of the new tool into a theory of mind.

The New Mind

What was later called the “new mental chemistry” pictured the mind as a computer program:

The atoms of this mental chemistry are symbols, which are combinable into larger and more complex associational structures called lists and list structures. The fundamental “reactions” of the mental chemistry employ elementary information processes that operate upon symbols and symbol structures: copying symbols, storing symbols, retrieving symbols, inputting and outputting symbols, and comparing symbols. (Simon, 1979, p. 363)

This atomic view is certainly a major conceptual change in the views about problem solving compared to the theories of Köhler, Wertheimer, and Duncker, but it bears much resemblance to the combinatorial view of intelligence of the Enlightenment philosophers.⁵

The different physical levels of a computer lead to Newell’s cognitive hierarchy, which separates the knowledge-level, symbol-level, and register-transfer levels of cognition. The seriality of 1971-style computers is actually embedded in Newell’s cognitive theory (Arbib, 1993).

One of the major concepts in computer programming that made its way into the new models of the mind is the decomposition of complexity into simpler units, such as the decomposition of a program into a hierarchy of simpler subroutines or into a set of production rules. On this analogy, the most complex processes in psychology, such as scientific discovery, can be explained through simple subprocesses. Thus, the possibility of the logic of scientific discovery, the existence of which Karl Popper so vehemently disclaimed, has returned in the analogy between computer and mind (Langley, Simon, Bradshaw, & Zytkow, 1987).

The first general statement of Newell and Simon’s new vision of mind appeared in their 1972 book, *Human Problem Solving*. Newell and Simon argued for the idea that higher level cognition proceeds much like the behavior of a production system—a formalism from computer science (and before that symbolic logic) that had never before been used in psychological modeling. Newell and Simon (1972) wrote of the influence of programming concepts on their models:

Throughout the book we have made use of a wide range of organizational techniques known to the programming world: explicit flow control, subroutines, recursion, iteration statements, local naming, production systems, interpreters, and so on. . . . We confess to a strong premonition that the actual organization of human programs closely resembles the production system organization. (Newell & Simon, 1972, p. 803)

Here we do not attempt to probe the depths of how Newell and Simon’s ideas of information processing changed theories of mind; the commonplace usage of computer terminology in the cognitive psychological literature since 1972 is a reflection of this. How natural it seems for present-day psychologists to speak of cognition in terms of *encoding, storage, retrieval, executive processes, algorithms, and computational cost*.

⁵ In fact, the new view was directly inspired by 19th-century mathematician George Boole (1854/1958), who, in the very spirit of the Enlightenment mathematicians such as the Bernoullis and Laplace, set out to derive the laws of logic, algebra, and probability from what he believed to be the laws of human thought. Boole’s algebra culminated in Whitehead and Russell’s (1935) *Principia Mathematica*, describing the relation between mathematics and logic, and in Claude Shannon’s seminal work (his master’s thesis at MIT in 1937), which used Boolean algebra to describe the behavior of relay and switching circuits (McCorduck, 1979, p. 41).

New Experiments, New Data

The tools-to-theories heuristic implies that new theories need not be a consequence of new experiments and new data. Furthermore, we argue that new tools can transform the kinds of experiments performed and data collected. This consequence of the tools-to-theories heuristic is also known to have happened when statistical tools turned into theories of mind (and around the same time).

One such case is the revolution of psychophysics through a new tool called *Neyman-Pearsonian hypothesis testing* (Gigerenzer & Murray, 1987). The new theory of mind inspired by this tool is known as *signal-detection theory*. The Neyman-Pearson technique deals with two kinds of errors, Type I and Type II (or false alarms and misses). When Tanner and Swets (1954) projected the tool into the mind, stimulus detection and discrimination—earlier understood in terms of “thresholds”—then became seen as a decision between two competing hypotheses based on a criterion that balances the probability of two kinds of errors. Consequently, the avalanche of experiments on auditory and visual detection and discrimination that followed their proposal kept track of both kinds of error in subjects’ judgments. The important point is that earlier experiments, such as the classical works of Fechner and Thurstone, paid attention to only one kind of error (Gigerenzer, 1994). What happened is that the new statistical tool inspired a new theory of mind, which in turn changed the kind of data generated in research. In this way, Tanner and Swets were able, in good conscience, to discard the years of contradicting results that preceded them.

A similar story is to be told with the conceptual change brought about by Newell and Simon—it mandated a new type of experiment that in turn involved new kinds of subjects, data, and justification. In academic psychology of the day, the standard experimental design, modeled after the statistical methods of Ronald Fisher, involved many subjects and randomized treatment groups. The 1958 *Psychological Review* article used the same terminology of *design of the experiment* and *subject* but radically changed their meanings. There were no longer groups of human or animal subjects. There was only one subject—an inanimate being, Logic Theorist. There was no longer an experiment in which data are generated by either observation or measurement. Experiment took on the meaning of simulation.

In this new kind of experiment, the data were of an unforeseen type—computer printouts of the intermediate results of the program. These new data, in turn, required new methods of hypothesis testing. How did Newell and Simon tell if their program was doing what minds do? There were two methods. For Newell and Simon, simulation was a form of justification itself: A theory that is coded as a working computer program shows that the processes it describes are, at the very least, sufficient to perform the task, or, in the more succinct words of Simon (1992), “A running program is the moment of truth” (p. 155). Furthermore, a stronger test of the model is made by comparing the output of the computer to the think-aloud protocols of human subjects.

Although all of this was a methodological revolution in the experimental practice of the time, some important parallels exist between the new information-processing approach and the turn-of-the-century German approach to studying mental processes. These parallels concern the analysis of individual subjects (rather than group means), the use of think-aloud procedures, and the status of the subject. In the early German psychology, as well as in American psychology of the time (until about the 1930s), the unit of analysis was the individual person, not the average of a group (Danziger, 1990). The two most prominent kinds of data in early German psychology were reaction times and introspective reports. Introspective reports have been frowned upon ever since the inception of American behaviorism, but think-aloud protocols, their grandchildren, are back (as are reaction times). Furthermore, in the tradition of the Leipzig (Wundt) and Würzburg

(Külpe) schools, the subject was more prestigious and important than the experimenter. Under the assumption that the thought process is introspectively penetrable, the subject, not the experimenter, was assumed to provide the theoretical description of the thought process. In fact, the main experimental contribution of Külpe, the founder of the Würzburg school, was to serve as a subject, and it was often the subject who published the article. In the true spirit of these schools, Newell and Simon put their subject, the LT, as a co-author of a paper submitted to the *Journal of Symbolic Logic*. Regrettably, the paper was rejected (as it contained no new results from the point of view of modern logic), and the LT never tried to publish again.

Acceptance

The second dose of information processing administered to psychology (after Newell & Simon's, 1972, *Psychological Review* article) was the Research Training Institute on the Simulation of Cognitive Processes at the RAND Corporation, organized by Newell and Simon. At the institute, lectures and seminars were conducted; IPL-IV programming was taught; and the LT, the General Problem Solver, and the EPAM model of memory were demonstrated on the RAND computer. In attendance were some scientists who would eventually develop computer-simulation methods of their own—including George Miller, Robert Abelson, Bert Green, and Roger Shepard.

An early but deceptive harbinger of acceptance for the new information-processing theory was the publication, right after the summer institute, of *Plans and the Structure of Behavior* (Miller et al., 1960). Despite the aforementioned 1959 dispute with Newell and Simon over the ownership and validity of the ideas within, this book drew a good deal of attention from all of psychology.

It would seem the table was set for the new information-processing psychology; however, it did not take hold. Simon (1991, p. 232) complained of the psychological community, which took only a “cautious interest” in Newell and Simon's ideas. The “acceptance” part of the tools-to-theories thesis can explain this: Computers were not yet entrenched in the daily routine of psychologists, as we show here.

No Familiar Tools, No Acceptance

We take two institutions as case studies to demonstrate the part of the tools-to-theories hypothesis that concerns acceptance—the Harvard University Center for Cognitive Studies and Carnegie-Mellon University (CMU). The former never came to embrace fully the new information-processing psychology; the latter did but after a considerable delay. Tools-to-theories might explain both phenomena.

George Miller, the co-founder of the Center for Cognitive Studies, was certainly a proponent of the new information-processing psychology. As we said, Miller et al.'s (1960) *Plans and the Structure of Behavior* was so near to Newell et al.'s (1958) ideas that it was at first considered a form of theft, although the version of the book that did see the presses is filled with citations recognizing Newell et al. Given Miller's enthusiasm, one might expect the Center, partially under Miller's leadership, to blossom into information-processing research. It never did. Looking at the 1963–1969 annual reports (Harvard University Center for Cognitive Studies, 1963, 1964, 1966, 1968, 1969), we found only a few symposia or papers dealing with computer simulation.

Although the Center had a PDP-4C computer, and the reports anticipated the possibility of using it for cognitive simulation, as late as 1969 it never happened. The reports mention that the

computer served to run experiments, demonstrate the feasibility of computer research, and draw visitors to the laboratory. However, difficulties involved in using the tool were considerable. The PDP saw 83 hr of use on an average week in 1965–1966, but 56 of these were spent on debugging and maintenance. In the annual reports are several remarks of the type, “It is difficult to program computers. . . . Getting a program to work may take months.” The Center even turned out a 1966 technical report entitled *Programmanship, or How to Be One-Up on a Computer Without Actually Ripping Out Its Wires*.

What might have kept the Harvard computer from becoming a metaphor of the mind was that the researchers could not integrate this tool into their everyday laboratory routine. The tool even turned out to be a steady source of frustration. As tools-to-theories suggests, this lack of entrenchment in everyday practice accounted for the lack of acceptance of the new information-processing psychology. Simon (1979) took notice of this:

Perhaps the most important factors that impeded the diffusion of the new ideas, however, were the unfamiliarity of psychologists with computers and the unavailability on most campuses of machines and associated software (list processing programming languages) that were well adapted to cognitive simulation. The 1958 RAND Summer Workshop, mentioned earlier, and similar workshops held in 1962 and 1963, did a good deal to solve the first problem for the 50 or 60 psychologists who participated in them; but workshop members often returned to their home campuses to find their local computing facilities ill-adapted to their needs. (p. 365)

At CMU, Newell, Simon, a new information-processing-enthusiastic department head, and a very large National Institute of Mental Health (NIMH) grant were pushing “the new IP [information processing] religion” (Simon, 1994). Even this concerted effort failed to proselytize the majority of researchers within their own department. This again indicates that entrenchment of the new tool in everyday practice was an important precondition for the spread of the metaphor of the mind as a computer.

Acceptance of Theory Follows Familiarity With Tool

At CMU in the late 1950s, the first doctoral theses involving computer simulation of cognitive processes were being written (H. A. Simon, personal communication, June 22, 1994). But this was not representative of the national state of affairs. In the mid-1960s, a small number of psychological laboratories were built around computers, including those of CMU, Harvard, Michigan, Indiana, MIT, and Stanford (Aaronson, Grupsmith, & Aaronson, 1976, p. 130). As indicated by the funding history of NIMH grants for cognitive research, the amount of computer-using research tripled over the next decade. In 1967, only 15% of the grants being funded had budget items related to computers (e.g., programmer salaries, hardware, supplies); by 1975, this figure had increased to 46%. The late 1960s saw a turn toward mainframe computers that lasted until the late 1970s, when the microcomputer started its invasion of the laboratory. In the 1978 Behavioral Research Methods & Instrumentation conference, microcomputers were the issue of the day (Castellan, 1981, p. 93). By 1984, the journal *Behavioral Research Methods & Instrumentation* appended the word *Computers* to its title to reflect the broad interest in the new tool. By 1980, the cost of computers had dropped an order of magnitude from what it was in 1970 (Castellan, 1981, 1991). During the last two decades, computers have become the indispensable research tool of the psychologist.

After the tool became entrenched in everyday laboratory routine, a broad acceptance of the view of the mind as a computer followed. In the early 1970s, information-processing psychology finally caught on at CMU. Every CMU-authored article in the proceedings of the 1973 Carnegie

Symposium on Cognition mentions some sort of computer simulation. For the rest of the psychological community, which was not as familiar with the tool, the date of broad acceptance was years later. Simon (1979) estimated that, from about 1973 to 1979, the number of active research scientists working in the information-processing vein had “probably doubled or tripled” (p. 390).

This does not mean that the associated methodology became accepted as well. It clashed too strongly with the methodological ritual that was institutionalized during the 1940s and 1950s in experimental psychology. We use the term ritual here for the mechanical practice of a curious mishmash between Fisher’s and Neyman-Pearson’s statistical techniques, which was taught to psychologists as the *sine qua non* of scientific method (Gigerenzer, 1993). Most psychologists assumed, as the textbooks told them, that there is only one way to do good science. But their own heroes—Fechner, Wundt, Pavlov, Köhler, Bartlett, Piaget, Skinner, and Luce, to name a few—had never used this “ritual.” Some had used experimental practices that resembled the newly proposed methods used to study the mind as computer.

Pragmatics

Some of our experimental colleagues have objected to our earlier analysis of how statistical tools turned into theories of mind. They have argued that tools are irrelevant in discovery and that our tools-to-theories examples are merely illustrations of psychologists being quick to realize that the mathematical structure of a tool (e.g., ANOVA) is precisely that of the mind. It is not easy to convince someone who believes (in good Neoplatonic fashion) that today’s theory of mind exactly fits the nature of the mind—that such a splendid theory might mirror something other than reality pure and simple. If it were true that tools have no role in discovery and that the new theories just happen to mirror the mathematical structure of the tool, then the pragmatics of the use of a tool—which is independent of the mathematical structure—would find no place in the new theories. In this section, however, we provide evidence that not only the new tool but also its pragmatic uses are projected into the mind. The tools-to-theories heuristic cannot be used to defend a spurious Neoplatonism.

One example is Kelley’s (1967) causal attribution theory, which postulates that the mind draws a causal inference in the same way social scientists do, by using Fisher’s ANOVA. The pragmatics, in addition to the mathematics of ANOVA, were projected into the mind. The practical use of a tool is generally undetermined by its mathematical structure. The mathematics of significance testing, as in ANOVA, has been used both for rejecting hypotheses based on data and for rejecting data (e.g., outliers in astronomical observations) based on hypotheses. Scientists have to get rid of both bad hypotheses and bad data. In the psychological laboratories, however, ANOVA was (and is) used almost exclusively for rejecting hypotheses based on data. Dubious data, in contrast, were (and still are) dealt with informally. When Kelley projected ANOVA into the mind, this specific, practical use (i.e., rejecting hypotheses) was projected along with it (Gigerenzer, 1991). In sharp contrast to earlier theoretical accounts, such as Michotte’s and Piaget’s, causal inference was seen as data driven, as an inductive inference from data to causes. Kelley’s new mind used the tool in the same way the researcher uses the tool—to trust the data (the information given) and to mistrust the hypotheses. The inductive view of causal attribution became one of the classic topics of social psychology, even to the point of defining the field.

The same process of projecting pragmatic aspects of the use of a tool into a theory can be shown for the view of the mind as a computer. One example is Levelt’s (1989) model of speaking. The basic unit in Levelt’s model, which he called the “processing component,” corresponds to the

computer programmer's concept of a subroutine. We argue that Levelt's model not only borrowed the subroutine as a tool but also borrowed the practical aspects of how subroutines are used and constructed in computer programming.

A subroutine (or "subprocess") is a group of computer instructions (usually serving a specific function) that are separated from the main routine of a computer program. It is common for subroutines to perform often needed functions, such as extracting cube roots or rounding numbers. There is a major pragmatic issue involved in writing subroutines that centers on the "principle of isolation" (Simon & Newell, 1986). The issue is whether subroutines should be black boxes or not. According to the principle of isolation, the internal workings of the subroutine should remain a mystery to the main program, and the outside program should remain a mystery to the subroutine. Black-box subroutines have become known as *program modules*, perfect for the divide-and-conquer strategy programmers often use to tackle large problems. To the computer, however, it makes no difference whether subroutines are isolated or not. Subroutines that are not isolated work just as well as those that are. The only real difference between the two types of subroutine is psychological. Subroutines that violate the principle of isolation are more difficult for the programmer to read, write, debug, maintain, and reuse. For this reason, introductory texts on computer programming stress the principle of isolation as the very essence of good programming style.

The principle of isolation—a pragmatic feature of using subroutines as a programming tool—has a central place in Levelt's model, in which the processing components are "black boxes" that exemplify Fodor's notion of informational encapsulation (Levelt, 1989, p. 15). In this way, Levelt's psychological model embodies a maxim of good computer programming—the principle of isolation. That this practical aspect of the use of the tool shaped a theory of speaking is not an evaluation of the quality of the theory. Our point concerns origins, not validity. However, this pragmatic feature of subroutines has not always served the model well. Kita (1993) and Levinson (1992) have attacked Levelt's model at its Achilles' heel—its insistence on isolated processing components.

To summarize the second part of this article, we started with the separation between intelligence and calculation and argued that the realization that computers can do more than arithmetic was an important precondition for the view of the mind as a computer. Newell and Simon seem to have been the first who tried to understand the mind in terms of a computer program, but the acceptance of their information-processing view was delayed until the psychologists became used to computers in their daily laboratory routine. We have argued that, along with the tool, its pragmatic use has been projected into theories of mind. Now that the metaphor is in place, many find it difficult to see how the mind could be anything else: To quote Philip Johnson-Laird (1983): "The computer is the last metaphor; it need never be supplanted." (p. 10)

Conclusion

Leibniz (1690/1951) once compared science to "an ocean, continuous everywhere and without a break or division" (p. 73). In this century, Hans Reichenbach (1938) parted this ocean into two great seas—the context of discovery and the context of justification—and Popper and his followers drew a sharp division between these two seas. Philosophers, logicians, and statisticians claimed justification as their territory and dismissed the context of discovery as none of their business and "irrelevant to the logical analysis of scientific knowledge" (Popper, 1959, p. 31). Popper relegated the study of discovery to the lower leagues of psychology and sociology: Either discovery is not worthy of serious study, or those who study it are not worthy of serious attention.

Thus, over one part of Leibniz's ocean shines the bright sun of philosophers, statisticians, and logicians, enlightening us about matters of justification; the other part, however, remains in mystical darkness, where mere intuition or chance reigns, or so it is claimed. However, there is a cry of protest among those who study science, insisting on the relevance of discovery and the promise of describing some of its general principles—something beyond the stock-in-trade anecdotes about Fechner, Kekulé, and Poincaré, which link discovery to beds, bicycles, and bathrooms. The challenge is to find general principles of discovery that go beyond the fascinating details of the individual genius.

The tools-to-theories heuristic is one of these principles. It connects the two seas and provides evidence that the commonly assumed fixed temporal order between discovery and justification—discovery first, justification second—is not necessary. We have discussed cases of discovery in which tools for justification came first and discovery followed (Gigerenzer, 1991; Gigerenzer & Murray, 1987). Moreover, in the present case study, we have drawn a full circle from theories of mind to computers and back to theories of mind. The argument was that economic changes—the large-scale division of labor in manufacturing and in the *bureaux de calculs*—went along with the breakdown of the Enlightenment conception of the mind, in which calculation was the distinctive essence of intelligence. After calculation was separated from the rest of intelligence and relegated to the status of a dull and repetitive task, Babbage could envision replacing human computers with mechanical ones. Both human and mechanical computers manufactured numbers as the factories of the day manufactured goods. In the 20th century, the technology became available to make Babbage's dream a reality. Computers became indispensable scientific tools for everything from number crunching to simulation. Our focus was on the work by Simon and Newell and their colleagues, who proposed the tool as a theory of mind. Their proposal reunited mere calculation with what was now called "symbol processing," returning to the Enlightenment conception of mind. After computers found a place in almost every psychological laboratory, broad acceptance of the metaphor of the mind as computer followed.⁶

Where do new ideas come from? Our argument is that discovery can be understood as more than mere inductive generalization or Popper's lucky guesses. More than that, we argue that, for a group of cognitive theories, neither induction from data nor lucky guesses played an important role. Rather, we propose that these innovations can be accounted for by the tools-to-theories heuristic. Tools have the power to inspire new theories. Even the mind has been re-created in their image.

Acknowledgements

We thank Peter Barker, Lorraine Daston, Berna Eden, Ralph Hertwig, Johannes Hönekopp, Joel Goldstein, Howard Gruber, David Klahr, David McNeill, Arthur Miller, Frank Renkewitz, Peter Sedlmeier, Herbert Simon, and William Wimsatt for their many helpful comments.

⁶ Our reconstruction of the path "from mind to computer and back" also provides an explanation for one widespread type of resistance against the computer metaphor of mind. The post-Enlightenment divorce between intelligence and calculation still holds to this day, and, for those who still associate the computer with mere calculation (as opposed to symbol processing), the mind-as-a-computer is a contradiction in itself.

References

- Aaronson, D., Grupsmith, E., & Aaronson, M. (1976). The impact of computers on cognitive psychology. *Behavioral Research Methods & Instrumentation*, 8, 129–138.
- Arbib, M. A. (1993). Allen Newell, unified theories of cognition. *Artificial Intelligence*, 59, 265–283.
- Aspray, W. (1990). *John von Neumann and the origins of modern computing*. Cambridge, MA: MIT Press.
- Baars, B. J. (1986). *The cognitive revolution in psychology*. New York: Guilford.
- Babbage, C. (1994). *Passages from the life of a philosopher* (M. Campbell-Kelley, Ed.). Piscataway, NJ: IEEE Press. (Original work written c. 1812)
- Boole, G. (1958). *An investigation of the laws of thought*. New York: Dover. (Original work published 1854)
- Castellan, N. J. (1981). On-line computers in psychology: The last 10 years, the next 10 years—The challenge and the promise. *Behavioral Research Methods & Instrumentation*, 13, 91–96.
- Castellan, N. J. (1991). Computers and computing in psychology: Twenty years of progress and still a bright future. *Behavior Research Methods, Instruments, & Computers*, 23, 106–108.
- Colebrooke, H. (1825). Address on presenting the Gold Medal of the Astronomical Society to Charles Babbage. *Memoirs of the Astronomical Society*, 1, 509–512.
- Danziger, K. (1990). *Constructing the subject: Historical origins of psychological research*. Cambridge, UK: Cambridge University Press.
- Daston, L. (1988). *Classical probability in the Enlightenment*. Princeton, NJ: Princeton University Press.
- Daston, L. (1992). The naturalized female intellect. *Science in Context*, 5, 209–235.
- Daston, L. (1994). Enlightenment calculations. *Critical Inquiry*, 21, 182–202.
- Edgington, E. S. (1974). A new tabulation of statistical procedures used in APA journals. *American Psychologist*, 29, 25–26.
- Galison, P. (1987). *How experiments end*. Chicago: University of Chicago Press.
- Gigerenzer, G. (1991). From tools to theories: A heuristic of discovery in cognitive psychology. *Psychological Review*, 98, 254–267.
- Gigerenzer, G. (1992). Discovery in cognitive psychology: New tools inspire new theories. *Science in Context*, 5, 329–350.
- Gigerenzer, G. (1993). The superego, the ego, and the id in statistical reasoning. In G. Keren & G. Lewis (Eds.), *A handbook for data analysis in the behavioral sciences: Methodological issues* (pp. 311–319). Hillsdale, NJ: Erlbaum.
- Gigerenzer, G. (1994). Where do new ideas come from? In M. A. Boden (Ed.), *Dimensions of creativity* (pp. 53–74). Cambridge, MA: MIT Press.
- Gigerenzer, G., & Murray, D. J. (1987). *Cognition as intuitive statistics*. Hillsdale, NJ: Erlbaum.
- Gigerenzer, G., Swijtink, Z., Porter, T., Daston, L., Beatty, J., & Krüger, L. (1989). *The empire of chance: How probability changed science and everyday life*. Cambridge, UK: Cambridge University Press.
- Gleick, J. (1992). *Genius: The life and science of Richard Feynman*. New York: Pantheon.
- Gooding, D., Pinch, T., & Schaffer, S. (Eds.). (1989). *The uses of experiment: Studies in the natural sciences*. Cambridge, UK: Cambridge University Press.
- Hacking, I. (1983). *Representing and intervening*. Cambridge, UK: Cambridge University Press.
- Harvard University Center for Cognitive Studies. (1963). *Third annual report*. Cambridge, MA: Harvard University, Center for Cognitive Studies.
- Harvard University Center for Cognitive Studies. (1964). *Fourth annual report*. Cambridge, MA: Harvard University, Center for Cognitive Studies.
- Harvard University Center for Cognitive Studies. (1966). *Sixth annual report*. Cambridge, MA: Harvard University, Center for Cognitive Studies.
- Harvard University Center for Cognitive Studies. (1968). *Eighth annual report*. Cambridge, MA: Harvard University, Center for Cognitive Studies.
- Harvard University Center for Cognitive Studies. (1969). *Ninth annual report*. Cambridge, MA: Harvard University, Center for Cognitive Studies.
- Johnson-Laird, P. N. (1983). *Mental models*. Cambridge, UK: Cambridge University Press.
- Kelley, H. H. (1967). Attribution theory in social psychology. In D. Levine (Ed.), *Nebraska symposium on motivation* (pp. 192–238). Lincoln: University of Nebraska Press.
- Kelley, H. H., & Michaela, I. L. (1980). Attribution theory and research. *Annual Review of Psychology*, 31, 457–501.
- Kita, S. (1993). *Language and thought interface: A study of spontaneous gestures and Japanese mimetics*. Doctoral dissertation, University of Chicago.

- Langley, P. W., Simon, H. A., Bradshaw, G. L., & Zytkow, J. M. (1987). *Scientific discovery: Computational explorations of the creative processes*. Cambridge, MA: MIT Press.
- Leibniz, G. W. von. (1951). The horizon of human doctrine. In P. P. Wiener (Ed.), *Selections* (pp. 73–77). New York: Scribner's. (Original work published 1690)
- Lenoir, T. (1986). Models and instruments in the development of electrophysiology 1845–1912. *Historical Studies in the Physical and Biological Sciences*, 17, 1–54.
- Levelt, W. J. M. (1989). *Speaking: From intention to articulation*. Cambridge, MA: MIT Press.
- Levinson, S. (1992). *How to think in order to speak Tzeltal*. Unpublished manuscript, Max Planck Institute for Psycholinguistics, Cognitive Anthropology Group, Nijmegen, The Netherlands.
- McCorduck, P. (1979). *Machines who think*. San Francisco: Freeman.
- McCulloch, W. S. (1965). *Embodiments of mind*. Cambridge, MA: MIT Press.
- McCulloch, W. S., & Pitts, W. (1943). A logical calculus of the ideas immanent in nervous activity. *Bulletin of Mathematical Biophysics*, 5, 115–133.
- Michotte, A. (1963). *The perception of causality*. London: Methuen. (Original work published 1946)
- Miller, G. A., Galanter, E., & Pribram, K. H. (1960). *Plans and the structure of behavior*. New York: Holt, Rinehart & Winston.
- von Neumann, J. (1958). *The computer and the brain*. New Haven, CT: Yale University Press.
- Newell, A., Shaw, J. C., & Simon, H. A. (1958). Elements of a theory of human problem solving. *Psychological Review*, 65, 151–166.
- Newell, A., & Simon, H. A. (1972). *Human problem solving*. Englewood Cliffs, NJ: Prentice-Hall.
- Piaget, J. (1930). *The child's conception of causality*. London: Kegan Paul, Trench & Trubner.
- Popper, K. R. (1959). *The logic of scientific discovery*. New York: Basic.
- Reichenbach, H. (1938). *Experience and prediction*. Chicago: University of Chicago Press.
- Schaffer, S. (1992). *Disappearing acts: On Gigerenzer's "Where do new ideas come from?"* Unpublished manuscript.
- Schaffer, S. (1994). Babbage's intelligence: Calculating engines and the factory system. *Critical Inquiry*, 21, 203–227.
- Simon, H. A. (1969). *The sciences of the artificial*. Cambridge, MA: MIT Press.
- Simon, H. A. (1979). Information processing models of cognition. *Annual Review of Psychology*, 30, 363–396.
- Simon, H. A. (1991). *Models of my life*. New York: Basic.
- Simon, H. A. (1992). What is an "explanation" of behavior? *Psychological Science*, 3, 150–161.
- Simon, H. A., & Newell, A. (1986). Information Processing Language V on the IBM 650. *Annals of the History of Computing*, 8, 47–49.
- Tanner, W. P., Jr., & Swets, J. A. (1954). A decision-making theory of visual detection. *Psychological Review*, 61, 401–409.
- Turing, A. M. (1950). Computing machinery and intelligence. *Mind*, 59, 433–460.
- Turing, A. M. (1969). Intelligent machinery. In B. Meltzer & D. Michie (Eds.), *Machine intelligence* (Vol. 5, pp. 3–23). Edinburgh: Edinburgh University Press. (Original work written 1947)
- Whitehead, A. N., & Russell, B. (1935). *Principia mathematica* (2nd ed., Vol. 1). Cambridge, UK: Cambridge University Press.