

Herbert Simon's Computational Models of Scientific Discovery

Stephen Downes

University of Cincinnati

1. Introduction

Herbert Simon's work on scientific discovery deserves serious attention by philosophers of science for several reasons. First, Simon was an early advocate of rational scientific discovery, contra Popper and logical empiricist philosophers of science (Simon 1966). This proposal spurred on investigation of scientific discovery in philosophy of science, as philosophers used and developed Simon's notions of "problem solving" and "heuristics" in attempts to provide rational accounts of scientific discovery (See Nickles 1980a, Wimsatt 1980). Second, Simon promoted and developed many of the crucial techniques and methods used in cognitive science. One is the use of computers to model internal psychological processes, a technique central to his account of scientific discovery. Another is protocol analysis, the use of the verbal reports of experimental subjects in psychology to construct accounts of their psychological processes. Protocol analysis is given a detailed formulation by Simon (Simon and Ericsson 1984), and is modified for use in the study of scientific cognition in the paper on Krebs (Kulkarni and Simon 1988). Third, Simon introduces normative proposals for science based on his computational investigations of scientific discovery (See also Zytkow and Simon 1988). Simon's work can be viewed as a contribution to naturalized philosophy of science, which centrally features the derivation of normative proposals from descriptive accounts of science. His work can also be viewed as a contribution to the growing field of "cognitive science of science," which uses techniques from the cognitive sciences to tackle issues in the philosophy of science.

In this paper I critically evaluate Simon's recent work on scientific discovery.¹ I focus primarily on *Scientific Discovery* (Langley et al. 1987), which documents many computer programs that purportedly make scientific discoveries, and "The Process of Scientific Discovery" (Kulkarni and Simon 1988), which is a detailed investigation of Krebs discovery of the ornithine cycle. I present several distinct criticisms of Simon's work. First I argue that Simon's descriptive account cannot distinguish discoveries as the product of an individual scientist's psychological processes or as the product of a social process involving several scientists. Yet Simon argues that scientific discovery can be adequately accounted for by appealing to psychological processes. I offer two lines of argument to establish that this latter step is unjustified. The first is that

Simon's method of protocol analysis does not provide sufficient evidence for the existence of the distinct psychological processes he claims underlie scientific discoveries. The second is that scientific discovery has a crucial social component that Simon cannot account for. I conclude that as a result of these failures Simon's descriptive characterization of scientific discovery is inadequate, and further that this inadequacy is due to what I call "cognitive individualism." I conclude the paper by considering the normative dimension of Simon's account, and argue that his computer models of scientific discovery can be best understood as contributions to what Clark Glymour has called "android epistemology."

2. An Outline of Simon's Project

In *Scientific Discovery* (Langley et al. 1987) Simon describes a set of computer systems that purportedly make scientific discoveries. Primarily he is concerned with systems that can make data driven discoveries, or discoveries of a particular relationship in a certain (usually numerical) data set. The systems use a set of production rules (conditional statements) that represent heuristics (rules of thumb), which constrain operations on the data. Such a rule (or heuristic) for detecting regularities in numerical data is: "If the values of two numerical terms increase together, then consider their ratio" (Langley et al. 1987, p.66). The "discovery" of Kepler's Third Law by such a system is characterized as a problem solving task, involving a search through the data, the "problem space," to produce the desired goal, the law. The data are values for the periods of the planets (P) and their distances from the sun (D), and the desired relationship (goal state) is $D^3/P^2 = c$.

Simon claims that the systems' potential application extends beyond the particular scientific discoveries investigated, to providing an account of scientific discovery in general. The computer programs, the BACON programs, and other related programs, are progressively developed to deal with gradually more complex data, for example data requiring qualitative laws and attribute ascription. *Scientific Discovery* documents a large array of programs that Simon claims have discovered Kepler's Third Law, Boyle's Law, Snell's Law and many more.

Simon claims that the work

...seeks to investigate the psychology of the discovery process, and to provide an empirically tested theory of the information-processing mechanisms that are implicated in that process. (Langley et al. 1987, p.4)

The work is firmly embedded in his overall information processing approach in cognitive science (See e.g. Simon 1969, Newell and Simon 1972). According to Simon, "the research is mainly limited to finding a set of mechanisms that is sufficient to account for discovery" (Langley et al. 1987, p.4). So Simon's claims are that scientific discovery is a psychological process, and that a sufficient account of this process will be provided by the computer models.

In "The Processes of Scientific Discovery: The Strategy of Experimentation" (Kulkarni and Simon 1988) Simon discusses Krebs experiments that led to the discovery of the ornithine cycle. Simon explains that the BACON programs did not broach the issue of where data came from in data driven discovery. He claims that the processes of designing experiments and observation were not investigated and that these latter are investigated in the work on Krebs. Simon's computational analysis of scientific discovery is governed by a guiding principle that scientific discovery is a collection of psychological processes and that these can be elicited by studying the work of

scientists throughout history who have made important discoveries. Therefore the production system resulting from the work on Krebs is intended to complement the BACON discovery systems.

Simon argues that he can provide a reliable account of the psychological processes involved in Krebs' discovery of the ornithine cycle by relying on Holmes' (Holmes 1980) historical work on Krebs'. Holmes' account is constructed from a combination of Krebs' published work, his and his assistant's laboratory notebooks and interviews with Krebs made years after the discovery. Holmes' account provides the necessary protocols from which Simon derives heuristics, the proposed psychological mechanisms that produced Krebs' discovery. These heuristics have been embodied in the production system KEKADA, which simulates Krebs' discovery. Overall Simon claims that Krebs' discovery of the ornithine cycle was due to a set of psychological processes that Krebs possessed, and that these are derived from his scientific writings used as protocols, and captured in the computer program KEKADA.

3. The Descriptive Adequacy of Simon's Account

Philosophers of science have turned to studies of scientific discovery as a reaction to the claims of Popper and the logical empiricist philosophers that discovery was not amenable to rational explanation (See e.g. Nickles 1980b). Scientific discovery is also the subject of investigation of historians and sociologists of science (See e.g. Brannigan 1981, Pickering 1984, Galison 1987). The picture of discovery that arises from these investigations is by no means monolithic, rather one of a complex and varied activity. Discovery is part psychological activity, part sociology of group acceptance, and part historical accident and timeliness. In contrast we see that Simon's account centers around the development of computer programs that arrive at the same results as great scientists in history. Simon argues that this approach provides a sufficient set of mechanisms to account for scientific discovery (Langley et al., p.4). Further he argues that the success of these computer programs will "show how simple information processes ... can give an adequate account of the discovery process" (Langley et al., p.33). Simon's descriptive account of scientific discovery shares none of the richness of the picture that arises from research in philosophy, sociology and history of science, and I argue that this makes it an insufficient account of scientific discovery.

Let us begin with the use of protocols. Simon's own observations about the weaknesses of the method of protocol analysis can be extended into arguments against his use of this method to elicit the psychological processes involved in scientific discovery. I conclude that protocol analysis provides insufficient evidence to support the existence of the proposed psychological processes underlying in scientific discovery.

Two methods of direct verbalizations are used to generate verbal reports from psychology subjects: Thinking aloud and retrospective accounts (My account follows Simon and Ericsson 1984 and Ericsson and Oliver 1988). The former are recorded as the subject carries out the task under study, and the latter are recorded immediately after the activity under study has taken place to make sure that the subject still has the relevant information in short term memory. In the paper on Krebs Simon claims that scientist's laboratory notebooks are closer in nature to retrospective reports than thinking aloud reports (Kulkarni and Simon 1988). We have no thinking aloud reports in the Krebs case as the protocols are from a historical study not a psychology experiment. When retrospective reports are used in psychology experiments they are made according to specific guidelines for remembering what was thought about during task performance (Ericsson and Oliver 1988). So Simon's comparison between laboratory notebooks and retrospective reports is a weak one. Although laboratory notes are

taken at the end of particular tasks (and even at the end of the day or the week), they are not taken at the end of a particular psychological process. The notion of coming to the end of a psychological process is not a relevant factor for a scientist determining when to make notes in their laboratory notebook.

Scientists do not primarily aim at recording their thought processes during experiments when making laboratory notebooks. For example Millikan's notebooks (see Holton 1978) contained columns of figures and comments such as "beautiful, publish this." This is good evidence that he was not always concerned with recording his psychological processes, rather with recording his results and commenting on their usefulness. Scientists in the laboratory record important results, or outline replicable procedures for themselves to use on a future occasion, or for graduate students or technicians to use in their absence.² Laboratory notebooks do provide useful data about scientific practice, but it is not necessarily evidence for the existence of particular psychological processes of a particular scientist.³ Such writings could also be used to generate an account of the social processes involved in a discovery (see Latour and Woolgar 1979). On the evidence of laboratory notebooks it is not only difficult to distinguish between different psychological processes, but it is also difficult to distinguish between the psychological processes of individual scientists and the more interactive processes of all the participants in the laboratory. The type of evidence provided by scientists' writings does not force one to the conclusion that particular psychological processes underlie scientific discovery.

In conclusion the analogy between verbal reports in psychology experiments, and laboratory notebooks as a resource for work on scientific discovery is strained for two reasons. First laboratory notebooks contain recollections, which may not have been made immediately after the putative psychological process they relate to occurred. Second, the laboratory notebooks do not contain information specifically about psychological processes. The laboratory notebooks can only be used as data from which an attempt to derive an account of psychological processes is made, and they could equally be used to derive an account of social processes.

The use of scientists more public writings to gain information about their psychological processes is even more problematic. One of the problems with retrospective reports in psychology is that subjects often "fill in" their reports with information not directly reproduced from memory (Simon and Ericsson 1984). They will perhaps give some plausible reasons for a particular activity instead of trying to remember the actual processes they went through (Cf. Nisbett and Wilson 1977). In a scientist's published work she is almost exclusively concerned with giving a plausible account of her results (in the form of reasons for these results), she is rarely if ever required to remember her psychological processes at the time of producing such results. Further she does not attempt to distinguish between reasons for results, and the psychological processes that led to such results. No principled method is available to distinguish between the two types of *post facto* reports: Ones that present psychological processes, and ones that present plausible reasons for a particular act (Cf. Nisbett and Wilson).

Leaving the issues surrounding protocols as insufficient evidence for the existence of psychological processes, a second line of argument challenges Simon's claim that such processes provide a "sufficient" account of scientific discovery. Drawing evidence from work in the sociology of science, I conclude that his cognitive individualist account of scientific discovery is not sufficient as it cannot account for the social nature of scientific discovery.

Simon's approach to explaining scientific discovery is directed by his information processing psychology. Simon investigates scientific discovery as a process of "thinking man" (or thinking machine) (Simon 1969). His position is that the scientists under investigation use one or more of a common stock of psychological processes. These processes are heuristic driven search mechanisms. An important question is why one person's use of some shared psychological processes would produce a scientific discovery, whilst another person's use of it would not. Simon claims in earlier work that the environment is the essential governing factor in producing different results with the same processes (Simon 1969, 1957), and yet he devotes no time in *Scientific Discovery* to describing *how* the environment is instrumental in producing the particular scientific discoveries he investigates. It is consistent with Simon's information processing psychology to argue that, given our shared psychological processes any human could come up with scientific discoveries, if he or she were put in the right environment. Putting it another way, it is consistent with Simon's account that factors other than the simple psychological processes he claims we possess are instrumental in producing scientific discoveries, and yet he leaves no room for such factors. Despite this deficiency he claims that he provides a sufficient account of scientific discovery. I now consider some factors that a sufficient account of scientific discovery must to account for.

Simon pays no attention to the issue of how one establishes that a scientific discovery has been made. A discovery's acceptance by the relevant scientific community is essential to its status as a discovery. And it is hard to separate this acceptance procedure from the process of discovery itself, a point argued by Brannigan in his *Social Basis of Scientific Discovery* (1981). Brannigan uses several examples from the history of science and exploring to illustrate his claim that the acceptance of a discovery by the relevant social group (the social context of the discovery), and the actual psychological process of discovery are indistinguishable (Cf. Woolgar 1988, pp. 58-65). For example he assesses Columbus' "discovery" of America (Brannigan 1981, pp. 120-142) and Mendel's "neglected" discovery of the genetic basis of inheritance (Brannigan 1981, pp. 89-119). In both cases Brannigan argues that the special social contexts determined these discoveries. He argues that it was the preparations for Columbus' voyage and the recognition of his achievement by royal sponsors that distinguish his discovery of America. And for Mendel it was the emergence of a context within modern biology for his work that rendered it a significant discovery, and it was not until such a context arose that it became a significant discovery.

The distinction between acceptance and discovery could perhaps be cashed out in terms of a clear distinction between cognitive and social factors of the scientific discovery process. On this distinction the cognitive component of the discovery would be that part explained by Simon's psychological models, say chronologically the part of a discovery up to the submission of a paper reporting the findings. The social component of the discovery could be the particular peer review process that led to the acceptance of the paper by a distinguished journal. But this hypothetical picture is too limited and obscures the complexity of scientific discovery. Simon's own account of Krebs' discovery of the ornithine cycle gives us enough information to question any account based on such a straightforward distinction between cognitive and social factors.

Krebs' work on the ornithine cycle was carried out with an assistant. For much of the time Krebs' assistant Henseleit did all the experimental work and took all the laboratory notes, whilst Krebs was pursuing more theoretical work on this and other projects (Holmes 1980). Simon paraphrases Holmes' account of the discovery, which includes an account of Henseleit's contributions, yet KEKADA models the putative psychological processes of an individual scientist. Whether Krebs *could* have carried

out the work leading to his discovery by himself is irrelevant here, as Simon aims to explain the *actual* discovery of the ornithine cycle (Kulkarni and Simon 1988, pp.140-143). But this discovery was produced by two cognitive agents whose interactions were instrumental in the discovery. KEKADA however is an idealized version of the possible psychological processes of an individual discoverer of the ornithine cycle. Thus KEKADA is not a model of the actual discovery of the ornithine cycle. Here we have a clear example of a scientific discovery, the relevant cognitive product, which was produced by more than one working scientist, or by social interaction.

If the interactive nature of scientific discovery were accepted, there would still be nothing in principle that prevents computer modelling of such activity.⁴ For example Simon could claim that he was modelling the discovery of the ornithine cycle by producing a production system that characterized two heuristic based problem solvers, and embodied them in a system that combined and synthesized their results. But Simon is a cognitive individualist with regard to scientific discovery.⁵ He holds that the cognitive process of scientific discovery can be accounted for by a model of an individual's psychological processes. The cognitive individualist approach prevents him from being able to provide a sufficient account of scientific discovery as it leaves important facets of scientific discovery unaccounted for, such as the interactions of a group of researchers essential to the eventual production of the scientific discovery, the relevant cognitive product.

For Simon scientific discoveries are produced by the psychological processes of an individual system, be it an individual scientist or computer program. Yet Simon's method of protocol analysis does not provide sufficient evidence to establish the existence of these psychological processes. Even if an account of the relevant psychological processes could be provided it would not provide a sufficient account of scientific discovery, as it cannot account for scientific discoveries arising from social interaction. Simon's cognitive individualist account cannot encompass the richness of scientific discovery revealed by sociologists and historians of science.

4. Normative Accounts of Scientific Discovery and Android Epistemology

Simon's computational models of discovery may not provide a sufficient account of scientific discovery, but perhaps they fulfill the role of refuting the claim, once held by the majority of philosophers of science, that scientific discovery is not amenable to rational analysis. Certainly Simon claims that his programs do this, but can this claim be sustained? Simon quotes what I call "Popper's challenge" from *The Logic of Scientific Discovery*:

The work of the scientist consists in putting forward and testing theories. The initial stage, the act of conceiving or inventing a theory, seems to me neither to call for logical analysis nor be susceptible to it (Langley et al. 1987, p.38).

In contrast Simon claims that there can be a normative theory of discovery. But if Simon has not provided a sufficient descriptive account of scientific discovery, what becomes of his normative account? I will suggest that Simon's normative account is best understood as a contribution to what Glymour has called "android epistemology" (Glymour 1987).

Logical empiricist philosophers and Popper concentrated on the development of theories of confirmation or justification of scientific theories. They assumed a split between discovery and justification; the former was not amenable to logical analysis, whilst the latter was. Since Hansson's work in the late fifties and sixties (see e.g. Hanson

1958), the most sustained philosophical discussion of scientific discovery in the literature is collected in Nickles' two volumes on scientific discovery (Nickles 1980b). The philosophers represented in this volume are still, like Popper and the logical empiricists, concerned with the normative dimension of science, but the "friends of scientific discovery" (Nickles 1980a) hold that normative accounts of scientific discovery are possible. For example some ways of going about scientific discovery are better than others. Most "friends of scientific discovery" have been concerned with elevating the status of scientific discovery from a mysterious process to a process amenable to rational analysis. The latter task involves providing a normative account of discovery. These philosophers also aim to provide accurate descriptive accounts of various scientific discoveries (See Nickles 1980b, Vol. II). They borrow historical or sociological methods to achieve this aim. An assumption driving this work is that if an accurate descriptive picture can be given of a great scientific discovery, it will inform the derivation of a normative account of scientific discovery in general (Nickles 1980a).

Simon has a similar project to the historically oriented philosophers of science (friends of scientific discovery), as he also aims to move from a descriptive to a normative account; the descriptive account constrains the norms derived. Simon focuses on both regulative and evaluative norms. Philosophers of science have traditionally been concerned with evaluative norms, for example norms for assessing a good theory. Regulative norms are those which, if followed, should produce effective procedures, including scientific discoveries. This distinction between regulative and evaluative norms is parallel to the one proposed by Nickles, who calls them "generative" and "consequentialist" norms (Nickles 1987).

Simon introduces his normative theory of discovery with the following claim: "The efficacy ("rationality," "logicality") [sic.] of the discovery process is as susceptible to evaluation and criticism as is the process of verification" (Langley et al. 1987 p.39). Simon is explicitly addressing Popper and the logical empiricists and their separation of the context of confirmation, or verification, from that of discovery. He claims that "a normative theory of discovery would be a set of criteria for judging the efficacy and the efficiency of processes used to discover scientific theories" (Langley et al., p.45). Simon claims that this theory "rests on contingent propositions such as 'If process X is to be efficacious for attaining goal Y, then it should have properties A,B, and C'" (ibid.) (the evaluative normative concern), and that "given such norms, we would be justified in saying that a person who adhered to them would be a better scientist" (ibid.) (the regulative normative concern).

Simon's work in *Scientific Discovery* is based on the assumption that there is no one scientific method, rather that there are several methods applicable over many domains of science. He calls these "weak methods" to contrast them with more powerful specific methods within a particular domain of research. According to Simon weak methods are to be judged against the limiting case of random search. He claims that scientists are very seldom involved in random search, and rational activity is distinguished from random search by the fact that the *best* use of weak methods is employed. So in the formula for the normative theory of discovery quoted above, variables A,B, and C correspond to weak methods. Simon goes on to substitute "heuristics" for the notion of "weak methods." Hence the normative theory of discovery is restated as: "Rationality for a scientist consists in using the best means he has available - the best heuristics - for narrowing the search down to manageable proportions" (Langley et al., 1987, p.47).

Simon's normative theory looks less like a normative theory of scientific discovery in general, than a theory of rationality construed as efficient search through a

problem space. Of course Simon's descriptive theory treats discovery as a form of problem solving, and on his account problem solving is just an heuristics based search through a problem space. So, from the point of view of his information processing perspective, rationality and rational scientific discovery may amount to nothing more than efficient search. But this is not a general normative theory of scientific discovery. It is still an open question whether Simon's descriptive account captures scientists' psychological processes, and hence whether a scientist who adopted Simon's regulative norms, or more specifically used the best heuristics, would make better discoveries. It may well be true that if scientists were information processors whose work was best characterized by search through a problem space, they would become better discoverers if they used the best heuristics available. But as we have already seen Simon's descriptive account of discovery is far too limited, so a normative theory derived from this account can have only a limited application.

Simon's work is minimally consistent with that of naturalistic philosophers of science (for example historically oriented philosophers of science) who claim that a normative account of scientific discovery can only be developed on the basis of, and at the same time as a descriptive account (Cf. Laudan 1977). The problem for Simon is that his descriptive account does not do justice to the complexity of the scientific discovery process, for example the social relations involved in the process. Consequently the normative account he derives can provide no directives for groups of scientists. Further it provides no directives for different instances of complexity in scientific practice, such as when it would be best to move to a different level of explanation to solve a scientific problem. In biology, for example, the solution to a particular problem might require shifting from the cellular to the biochemical level. What Simon's account does provide is a set of norms for guiding the simulation of further scientific discoveries, provided simulation can be achieved by representing the activity in terms of an heuristics based search through a problem space. Simon even claims his approach cannot "replicate the historical details of various scientific discoveries" (Langley et al., p.62). Instead it can provide models of how such discoveries "*might* occur." If this is the overall claim of *Scientific Discovery*, then it is an entirely normative one, but we have seen that Simon also aimed to provide a descriptive account. I will conclude by suggesting that Simon's normative claims may be best understood as guidelines to assist in building good scientific discovery machines, and so his work is best understood as a contribution to what Glymour has called "android epistemology" (Glymour 1987).

Glymour has proposed the name "android epistemology" for the production of the norms that regulate machines such as Buchanan and Mitchell's meta-DENDRAL (Buchanan and Mitchell 1978), which have attained some level of success in scientific discovery (Glymour 1987). This project involves the development of machines to solve problems that humans have been accustomed to solving, especially problems that have traditionally interested philosophers. Scientific discovery is of central interest to many philosophers, so for Glymour the development of norms for machines that make scientific discoveries is work in "android epistemology."

Glymour has proposed a logic of scientific discovery implemented as a computer program (Glymour et al. 1988, and Glymour and Kelly 1989). The explicit difference between Glymour's approach and Simon's is that Glymour is not concerned with modelling actual human psychological processes (Glymour 1987, 1988). Glymour claims that "conventional artificial intelligence programs are little theories. The more theories look like theories of reasoning, the more the description of the program looks like a piece of philosophy" (Glymour 1988, p.200). Glymour's view of philosophy is close to logical empiricism. He claims that the philosopher's concern is to give a "re-

construction of a domain of knowledge or form of reasoning" (Glymour 1988, p.201. Cf. Glymour 1981). Glymour shares with the logical empiricists the assumption that reasoning processes can be abstracted from their context. In the case of scientific practice it is argued that the scientists' reasoning can be appraised independently of other scientific practices. Glymour diverges from the logical empiricists in arguing that one *can* provide a logic of scientific discovery, which for him is a theory of the reasoning that produces scientific discoveries. He claims that AI programs provide him with the formal capability of presenting such a theory. Finally, Glymour's account is an entirely normative one. His normative theory of discovery is proposed as a theory of how to make the best scientific discoveries (regulative norms). Or even stronger: How to go about discovering the truth. Recently he has claimed that he can give an account of the reasoning involved in discovering "the truth and nothing but the truth" (Glymour and Kelly 1989). Setting the external standards by which his norms are judged very high.

Simon's and Glymour's normative accounts share the same goals. Simon aims to provide an account of how the best scientific discoveries will be made. The way the account is implemented is in the construction of computer programs. Such programs bear little relation to the actual practice of scientific discovery, and so they fulfill one of Glymour's requirements for contributions to android epistemology, as they do not replicate human endeavor (Glymour 1987). The requirement is that the android epistemologist avoid what Glymour calls the "anthropocentric constraint": "[T]hat the algorithms executed by an android in performing a task must, at some appropriate level of description, be the *very same* algorithms that people execute in performing that task" (Glymour 1987, p.74). While Glymour explicitly avoids the "anthropocentric constraint," Simon avoids it by default, due to the insufficiency of his descriptive account of scientific discovery.

Simon's and Glymour's accounts of scientific discovery are in direct competition if they are both understood as android epistemology. The decision between the two accounts may be determined by generality of application, say by the scope of the regulative norms in each account. If we consider the applicability of regulative norms in terms of the maxim "ought implies can," then neither account is generally applicable to human scientists. If one is concerned with the production of good science *per se*, then it is an empirical question which account is more generally applicable. The answer will depend on the quality of the new scientific work that the relevant norm-guided computers produce in the future.

Android epistemology is only one facet of the enterprise of cognitive science of science, others concern the description and explanation of human cognitive practices.⁶ Glymour explicitly rejects this latter project. Simon, on the other hand, arrives at android epistemology by default, due to the inadequacy of his descriptive account. He provides guidelines for the design of efficient computer programs. Simon's account of scientific discovery does little to increase our understanding of scientific discoveries made by humans throughout history, and provides no useful regulative norms for groups of scientists practicing research currently. These are two of the goals that a cognitive science of science might achieve, and Simon's work fails to reach them.

Some more general proposals about naturalistic yet normative philosophy of science arise from this analysis. The first is that if a naturalized philosophy of science is to make any claims of descriptive adequacy, it must make use of empirical work in sociology and history of science as well as in cognitive science. The second is that if we are concerned with providing either regulative or evaluative norms for science, we should be concerned with their potential for application. The naturalistic turn in phi-

losophy adds a new challenge to those who produce norms, which is that the norms must be applicable by human agents within certain constrained situations. The difference between android epistemology and a more general normative naturalistic epistemology is that the latter should aim to apply to human scientists who have limited psychological capacities and must always act in social contexts.

Notes

¹Throughout the paper I will refer to Simon's collaborative work by Simon's name alone. The citations credit his co-workers.

²It is worth noting that graduate students and technicians often make entries in the laboratory notebooks. Certainly they cannot be recording the internal psychological processes of their supervisor. We see below that Henseleit, Kreb's assistant, took many of the notes the Holmes study was based on, yet Simon treats these as protocols.

³Cf. Tweney (1989) who derives an interesting account of what he calls "external memory" from a study of Faraday's notebooks. (See also Gooding and James 1985.)

⁴Rob Cummins' SOFT program is an example of a program that models the cognitive activity of groups of people (Cummins 1983).

⁵It is important to note the qualification "with regard to scientific discovery." In Simon's work on administrative behavior and his use of the notion of satisficing, the claim that the individual was the prime unit of analysis was not central. See for example his *Models of Man* (1957), which is interestingly subtitled "Mathematical Essays on Rational Human behavior in a Social Setting." Simon's work in information processing psychology has many affinities with his work in organizational behavior, which was neutral with regards its units of analysis. The work was applicable to individuals or groups, such as business organizations. I detect a tension between Simon's work on organizations and his work on scientific discovery, the former is neutral over its units of analysis and the latter is cognitive individualist. A possible resolution would be to view scientific discoveries as produced by organizations, and so the relevant heuristics would govern group behavior. Simon nowhere indicates that this is the way his work on scientific discovery should be understood.

⁶Of course one cannot hold this sharp division between the goals of android epistemology and other more descriptive goals of cognitive science of science if one presupposes that humans and computers cognitive capacities are both computationally bounded. Much empirical work in cognitive science shows that although there are many deficiencies of human cognitive practices, the important bounds to human cognition are not purely computational (See eg. Faust 1984., cf. Cherniak 1986).

References.

- Baars, B.J. (1986), *The Cognitive Revolution in Psychology*. : New York: Guildford Press.
- Brannigan, A. (1981), *The Social Basis of Scientific Discoveries*. Cambridge, UK: Cambridge University Press.

- Buchanan, B.G. and Mitchell, T.M. (1978), "Model Directed Learning of Production Rules," in Waterman and Hayes-Roth (eds.) (1978).
- Cherniak, C. (1986), *Minimal Rationality*. Cambridge, Mass: MIT Press.
- R.G. Colodny, (ed.) (1966), *Mind and Cosmos*. Pittsburgh: University of Pittsburg Press.
- Cummins, R.C. (1983), "SOFT," in *The Proceedings of the Conference on Artificial Intelligence*, Oakland University.
- Ericsson, K.A. and Oliver, W.L. (1988), "Methodology for Laboratory Research on Thinking: Task Selection, Collection of Observations, and Data Analysis," in Sternberg and Smith (eds.), pp. 392-428.
- Faust, D. (1984), *The Limits of Scientific Reasoning*. Minneapolis: University of Minnesota Press.
- Fetzer, J.H. (ed.) (1988), *Aspects of Artificial Intelligence*. Dordrecht: Kluwer.
- Galison, P. (1987), *How Experiments End*. Chicago: University of Chicago Press.
- Gholson, B., Shadish, W.R., Neimeyer, R.A., and Houts, A.C. (eds.) (1989), *Psychology of Science: Contributions to Metascience*. Cambridge University Press: Cambridge, U.K.
- Glymour, C. (1988), "Philosophy is Artificial Intelligence," in Fetzer (ed.), pp. 195-207.
- (1987), "Android Epistemology and the Frame Problem," in Pylyshyn (ed.) (1987), pp. 65-75.
- (1981), *Theory and Evidence*. Princeton, NJ: Princeton University Press.
- and Kelly, K. (1989), "Convergence to the Truth and Nothing but the Truth," in *Philosophy of Science* 56.
- Glymour, C., Kelly, K. and Spirtes, P. (1988), "Philosophy of Science and the Logic of Discovery," unpublished manuscript.
- Gooding, D. and James, F.A. (eds.) (1985), *Faraday Reconsidered*. New York: Stockton Press.
- Hanson, N.R. (1958), *Patterns of Discovery*. Cambridge, U.K: University Press.
- Holton, G. (1978), *The Scientific Imagination*. Cambridge, U.K: Cambridge University Press.
- Kulkarni, D. and Simon, H.A. (1988), "The Process of Scientific Discovery: The Strategy of Experimentation," *Cognitive Science* 12: 139-176.
- Langley, P., Bradshaw, G.L, Simon, H.A. and Zyngow, J.M. (1987), *Scientific Discovery*. Cambridge, Mass: MIT Press.

- Latour, B. and Woolgar, S. (1979), *Laboratory Life*. London: Sage .
- Laudan, L. (1977), *Progress and its Problems*. Berkeley: University of California Press.
- Nersessian, N. (ed.) (1987), *The Process of Science*. Dordrecht: Nihjoff.
- Nickles, T. (1987), "Twixt Method and Madness," in Nersessian (ed.) 1987.
- (1980), "Scientific Discovery and the Future of Philosophy of Science," in Nickles (ed.), pp. 1-63.
- (ed.) (1980), *Scientific Discovery*. 2 Volumes, Dordrecht: Reidel.
- Nisbett, R. and Wilson, T.D. (1977) "Telling More than we can Know: Verbal Reports on Mental Processes," *Psychological Review* 84: 231-259.
- Pickering, A. (1984), *Constructing Quarks: A Sociological History of Particle Physics*. University of Chicago Press: Chicago.
- Pylyshyn, Z.W. (ed.) (1987), *The Robot's Dilemma*. Ablex: New Jersey.
- Simon, H.A. and Ericsson, A. (1984) *Protocol Analysis: Verbal Reports as Data*. MIT Press: Cambridge, Mass.
- (1977), *Models of Discovery*. Reidel: Dordrecht.
- and Newell, A. (1972) *Human Problem Solving*. Prentice-Hall: New Jersey.
- (1969), *The Sciences of the Artificial*. MIT Press: Cambridge, Mass.
- (1966), "The Psychology of Scientific Problem Solving," in , R.G.Colodny, (ed.) 1966.
- (1957), *Models of Man*. John Wiley and Sons, Inc.: New York.
- Sternberg, R.J. and Smith, E.E. Eds. (1988), *The Psychology of Human Thought*. Cambridge University Press: Cambridge, U.K.
- Tweney, R.D. (1989), "A Framework for the Cognitive Psychology of Science," in Gholson et al. (eds.), pp. 342-366.
- Waterman, D.A. and Hayes-Roth, F. (eds) (1978), *Pattern Directed Inference Systems*. Academic Press: New York.
- Wimsatt, W.C. (1980), "Reductionist Research Strategies and their Biases in the Units of Selection Controversy," in Nickles (ed.), 231-259.
- Woolgar, S. (1988), *Science: The Very Idea*. Tavistock: London.
- Zygtow, J.M. and Simon, H.A. (1988), "Normative Systems of Discovery and Logic of Search," *Synthese* 74: 65-90.