EDWIN HUTCHINS
University of California, San Diego

Response to Reviewers

The issues raised by the reviewers seem to revolve around two principal themes.

First, there is the question of the boundary of the unit of analysis and the implications of moving this boundary out beyond the individual. All of the reviewers agree that the new unit of analysis is a useful construct, but Latour feels that I am not consistent in it's application. Moving the boundary of the unit of analysis relocates some aspects of mind outside the individual. This raises three additional questions: Can all aspects of mind be delegated out? What remains of the mind of the individual when this process is complete? And finally, how shall we describe (or model) such distributed systems?

The second major theme concerns the setting in which the research was carried out. Ship navigation is a very specialized activity. Is it representative of many other settings for human activity? Can results from studies of settings like this be generalized? The reviewers contrast this setting with laboratory experiments on the one hand and freer, less historically stabilized human activity systems on the other. If ship navigation is different from other settings, how is it different? And what are the implications of these differences?

1. The Unit of Analysis

Latour is disturbed by my failure to "go all the way" with the idea of distribution. He says that when I write about "a person and the person's surroundings," I am slipping back into the way of thinking that I had hoped to challenge. This is surely a possibility, and the members of my laboratory and I have struggled for years to overcome the unwanted implications of familiar words. However, I'm also trying to explain to cognitive scientists how to see something familiar from a new perspective. I imagine an overhead transparency on which I have all the media of interest arranged in interaction with one another. There is no distinction between person and surroundings in this picture. This is a picture of a functioning distributed cognitive system. It consists only of media in coordination with one another. I have in my hand a separate transparency sheet on which I have drawn the old inside/outside boundary of a person. Note that this is the boundary of a person, not of a mind. I alternate between showing and describing the scene without the boundary in place and looking at the same scene with the boundary in place. In this way I show those who used to believe in putting all of cognition inside the person that cognition is bounded otherwise. I don't think this is "back-sliding," it's a deliberate device intended to permit viewers unaccustomed to the new perspective to learn how to see the world without the boundary drawn in so sharply.

I believe that Latour and I are in agreement about the need to dissolve the old boundaries and focus on the web of interacting media. But Latour seems to want to go further. He would dissolve the individual and the psychology of individuals as well. He would have psychology swept out and the person swept clean. He characterizes my thinking agent as "like the desk of a well-organized executive: empty since everything else has been delegated outside to something or someone else." I disagree. I would say that we have to reconceive much of the psychology that we have, but that the person is not left empty. This is precisely the point of chapter 7 of *Cognition in the Wild.* In that chapter I tried to erase the boundary of the skin and the skull while not erasing that which lay inside the old boundary. I tried to show how learning should be seen not as internalization (in the sense of the movement of something from outside to inside) but as the spread of organization in a complexly connected system. I don't want to empty the person - delegating all the work to someone or something else (as Latour would have it). If it's delegated to someone else, one wonders how that someone works. This other someone is presumably also empty, having delegated the delegated tasks elsewhere. This doesn't solve the problem, it only chases it out of sight. Rather, I would connect what is in the person to what is around the person. Now Latour will hit me again for having reconstituted the inside/outside boundary by talking about what is in and what is around. But it seems to me that the problem is not so much in acknowledging this boundary as it is in assuming that the boundary constitutes a clear separation of cognitive realms. This I do not do. Instead, I have attempted to develop a language of description of cognitive events that is unaffected by movement across the old boundaries. Latour recognizes the importance of this effort, and Keller correctly sees it as an attempt at "a uniform cognitive theory applicable to the diverse elements of activity systems." One cannot empty the person by delegating cognitive activity to "something or someone else." The work must be done somewhere, and some of the work will be done in regions that lie inside the bounds of persons.

Latour cites *Cognition in the Wild* in support of his claim for the final dissolution of psychology.

> The thinker in this world is a very special medium that can provide coordination among many structured media - some internal, some external, some embodied in artifacts, some in ideas, and some in social relationships.

This vague sentence points the way to the hard work of attempting to remake cognitive science. The really interesting question is this: What kind of medium is this thinker? The answer to this question is a new theory of psychological functioning. This is work that the book calls for but does not undertake. Is it possible to take this conception back to the long list of unanswered questions in cognitive science? My students and I have begun to revisit a small number of phenomena and processes that we believe can be better understood in terms of coordination of media and propagations of representational state. Our initial targets include selective attention and the phenomena of interpersonal coordination of attention, the acquisition of social skills, the structure and development of language, and the nature of expert performance. My guess is that when we are done, the mind will not be swept clean, but it will be furnished in a different way than it was in the past.

This brings us to a very difficult problem: How shall we describe (and model) such systems?

In my first book, I set myself the standard of AI of the time (the late 1970s): that a working program constitutes an explanation of a phenomenon. I produced a description of the problem, but not a program and I was rightly criticized for failing to meet my own standard.

Now, even though I no longer believe in the standard I held up for myself, and failed to meet, in my earlier book, I still do believe in the utility of computational modeling as a way of forcing one to be explicit about one's terms and claims. In *Cognition in the Wild,* I include some simple computational simulations of group processes. Latour refers to the simulations as "amusing." That is perhaps too kind, since they are schematic in the extreme. These are not explanations of any phenomena so much as they are existence proofs. They say: a system organized in this way could give rise to the sort of phenomena we are interested in. This is a weak sort of argument. It would be vastly preferable to be able to demonstrate how these things actually work.

Computational models of thought have changed dramatically in the past decade. And while all forms of computer program can be described in terms of the propagation of representational state, connectionist architectures would seem especially appropriate for this effort because they can be naturally thought of in terms of the coordination of media. But even the most advanced connectionist architectures do not support the rapid and flexible reconfigurations that seem to be characteristic of cognition in the wild. I see this lack of the right kind of computational framework is a major obstacle to progress.

All of this talk about formalism raises another tension in the reviews. From the point of view of these reviewers (and probably the readers of this journal), the book is certainly formal or rigorous enough. It is not, however, perceived in that way by much of the cognitive science audience I was trying to reach. For them it is squishy, and the ethnographic methods are unfamiliar and suspect.

2. The Setting

Let's turn now to the properties of the setting in which the work was carried out. Keller sees ship navigation as a real world activity which is to be contrasted with laboratory experiments. She sees in this approach a requirement for "a reorientation of traditional cognitive science research from laboratory tasks to real-life achievements." Bazerman sees ship navigation as "well-regulated" and "historically stabilized" and contrasts it with "collaborative work in education or industry, where the emphasis is more on individualized and improvisatory behaviors." Finally, Latour warns that some readers could see the quartermasters as ordinary people and could thereby fail to draw any lessons from this study for the work of "higher minds" such as scientists. Clearly, this setting could be different from other settings in an infinity of ways. Which ways matter, and why do they matter?

The theory claims that we should attend to the resources available to participants for use in organizing their behavior. That means that while different settings may consist of different sets of resources, they do not differ from each other in theoretically interesting ways.

Keller comments on a reorientation of traditional cognitive science research from laboratory tasks to real-life achievements. The issues here are complex. The main implication of the theory concerns how we look at cognition, not where we look for it. Of course, changing the way we look at cognition will have consequences for what we think can be learned from looking in various places, so that places that used to seem informative may subsequently appear less so.

It certainly seemed necessary to examine cognition in the wild in order to see the cultural nature of cognition. This is not because laboratory cognition is acultural, but because the tradition that pursues cognition in the laboratory has an investment in believing that the effects of culture can be controlled. Culture is not absent from the laboratory, but the setting is constructed in a way that directs attention away from it. So, while I do endorse a call for more studies of cognition in the wild, I want to be sure we are clear on why that is a good thing to do. The wrong reason for studying cognition in the wild is the belief that experiments are a special setting, or somehow different in cognitive or cultural terms from other real-world settings. The reorientation I advocate is not framed in terms of the settings of cognitive performances (lab versus real world) but in terms of the stance we take with respect to the interpretation of the observed behavior. As Lave, Suchman, and others have pointed out, laboratory experiments are just another socially organized context for performance. In a recent special issue of the journal *Cognitive Science,* a number of authors (mainly opponents of the situated action view) mistake situated action for action that takes place in real-world (i.e., non-laboratory) settings. In fact, cognition in the laboratory is just as situated as any other instance of cognition. The implication of the distributed cognition view is not that laboratory research should be abandoned in favor of "real-world" settings, but that the way that behavior that occurs in laboratory settings is interpreted should be changed to reflect the ways that subjects make use of cultural resources in the production of that behavior. Once this is done, however, some of the putative advantages of the laboratory disappear, and perhaps laboratory research becomes less interesting because it will seem to provide fewer answers than was assumed. To reiterate, we need to look in the wild, not because that is were the real cognition is, but because that is a place where it is easier to see the cultural nature of cognition.

Among real-world settings, is there anything that distinguishes ship navigation from other activities? First, while it is not unique, it is special in the sense that there is an easy to identify computation being performed by a distributed system. When I began, I had no idea how important this feature would be. The fact that many of the resources available to the participants are directly observable by the researcher does not change the theoretical standing of the setting, but it does make the analysis of the use of those resources much easier than it would otherwise be. Rationalized and historically stabilized settings where problems and their solutions have been crystallized in physical artifacts are simply easier to study than setting that lack that kind of structure. These are places where it is easier to see the role of cultural resources because there are plenty of them about for the participants to use. And the ones that are around are easy for the researcher to document and describe. When exploring a new theory, it is a good idea to tackle the methodologically easy cases first. Bazerman reminds us that the not so easy cases remain. Many human activities are difficult to characterize as computational in nature. This raises the question of the extent to which the approach I present here can be applied to other domains. I would like to believe that the problems will be mostly methodological, but I am prepared to discover new theoretical insights as we explore the range of applicability of this approach.

Bazerman is concerned with the rigid constraints of military life on the actions of the quartermasters. He notes that "Even when Hutchins observes a mechanical breakdown which requires improvisation of a new set of procedures, the improvised procedures rapidly move to a new set of regularities." Inserted as it is in a paragraph about the "narrowly defined pre-determined roles and behaviors" it gives the impression that this return to regularities is a property of the military context. I would argue, on the contrary, that the characterization of limited action roles as being typical of military organization is the product of a stereotype. The analysis of the development of a new set of regularities was intended as an example of how a very general cultural process works. The mechanisms that my analysis identified as responsible for the adaptation were not in any way grounded in "military discipline." They were, instead, ways of bringing psychological, social, and computational constraints into coordination with one another.

This aspect of the work raises much deeper questions about the origins and stability of structure. Is the stability I observed and the ability to adapt to perturbations a consequence of actors wanting the world to be regular and stable, or is it a consequence of the operation of a general set of adaptive processes operating in a world of constraints? In the most ambitious terms my goal would be to describe a system in which adaptive processes that are continually operating are responsible for the production of both stability and change.

Will this book change minds in cognitive science? Keller is optimistic. Bazerman is unsure. Latour is pessimistic. He says I am naive to think I can upset the psychology of the individual. I certainly don't expect this book by itself to do it. All of the reviewers see the ambition of the project clearly. I do intend this as a beginning on a project to remake cognitive science (and revitalize cognitive anthropology if enough of an ember remains of that field to be puffed back to flame). But there is an enormous amount of work yet to do. Cognitive scientists will remain unconvinced until there is a good computational model of the mechanisms involved and a clear demonstration that those mechanisms are not only capable of producing the phenomena of interest, but can do so better than the alternatives can. I don't expect to be the one to devise a computational scheme that is appropriate to express these ideas, but, if I can make the nature of the phenomena clear to others, perhaps someone will. The sort of theoretical stance adopted in *Cognition in the Wild* is obviously not the only challenge to the traditional symbolic approach. Of course this book alone will not change the field, and neither could it have been written if many of the relevant ideas were not already being developed by others with similar interests. I hope that this book and others like it might inspire a body of work that could change cognitive science. Is that naive? Perhaps. As Bazerman says, the outcome is in the hands of others.