Philosophical Psychology
Vol. 19, No. 1, February 2006, pp. 119–127

Reviews

Decisions, Uncertainty and the Brain: The Science of Neuroeconomics
PAUL GLIMCHER
Cambridge, MA: MIT Press, 2004
395 pages, ISBN: 0262572273 (pbk); $22.00

Most social scientists interested in neuroeconomics tell a story in which attention to psychological studies of rationality and decision making gave economics a new degree of psychological realism—both in itself and through the responses from mainstream economics that it stimulated. But these developments told us nothing about neuropsychological mechanisms until cognitive neuroscience developed a social dimension. Part of that development was the arrival of neuroeconomics, which studies the neural basis of economic behavior. The story concludes by looking forward to a synthesis linking economic theory—informed by psychology or sociology—with neuroscience in order to bring us an understanding of the behavior and biology of more or less rational agents.

Paul Glimcher came to neuroeconomics via oculomotor physiology, and he tells the story from another perspective, as the introduction of greater formal sophistication into neuroscience. Glimcher argues that the rational agent of neoclassical economics is in fact just what's needed to lead the way out of an impasse that has afflicted neuroscience since its inception. The key idea is that nervous systems are designed to maximize inclusive biological fitness via maximizing utility (p. 267). Glimcher's own work demonstrates the scientific progression that he foresees. While working on the visual system, he came to develop probabilistic approaches to investigate how monkeys choose to allocate attention. The book covers a lot of territory, and although philosophers of mind won't find much here to help with the traditional problems of their subject, there's a lot to think about for philosophers interested in how the cognitive neurosciences work, as well as other scholars who want to know more about neuroeconomics and its potential significance for cognitive psychology more broadly.

The problem Glimcher thinks neuroeconomics can address is that behavioral neuroscience has been driven by physiology—it has discovered brain circuits and then looked for their function. The program derives from Descartes' attempt to explain all human action as the reflexive response of the central nervous system to external stimuli. (Philosophers should pause sometimes to remember that Descartes was, for decades, known as one of the great champions of materialism.) This has been a bottom-up approach, in which reflex circuits are first understood as anatomical units and then investigated experimentally. The problem Glimcher sees with this

ISSN 0951-5089 (print)/ISSN 1465-394X (online)/06/01000119-127 © 2006 Taylor & Francis
DOI: 10.1080/09515080500462438

method is its failure to ask about the function of behaviors in the overall life of the organism. Glimcher draws on economics to suggest a rival top-down approach in which tools from probability and game theory are used to specify the tasks that an organism has evolved to perform. I'll try to tell the story the way Glimcher tells it, focusing on what I take to be the big picture. There's a lot of fascinating science in the book, but I will not go over it in detail.

To explain why contemporary issues have the form they do, Glimcher looks to their history; the exception, oddly, is neoclassical economics. The history is often agreeable to read—one gets the sensation of touring a well-stocked intellect—but the downsides are that crucial steps in Glimcher's argument keep getting deferred and some of the history seems irrelevant.

Still more is eccentric, or just hard to follow. For example, Glimcher claims that Sherrington and Pavlov wanted a mathematically complete model of the organism in which behavior was to be explained by understanding the relations between reflexes as akin to the relations between logical constants. This project, he says, was dealt a dire blow by Gödel's proof of the incompleteness of mathematics (pp. 72–73). But nothing more than first-order logic seems to be at stake in the discussion of Sherrington and Pavlov—i.e., if any logic is involved at all. None of the extensive quotations from Sherrington and Pavlov seem to support the claim that they were modeling reflex theory on mathematical logic. Glimcher does say that the physiologists just had in mind something similar to what the logicians were doing, but if the projects were just somehow similar then I don't understand the relevance of Hilbert and Gödel to the biology (p. 68). The limitations of mathematics are only significant if the reflex arc tradition was looking to do something that, mathematically, cannot be done.

This example shows what's wrong with the history, especially in the opening chapters. It lacks the detail to make Glimcher's historical case convincing. On the other hand, it has too much detail to merely be a summary of the development of the issues that Glimcher thinks cognitive neuroscience must now confront. The exception is his discussion of the experimental tradition in neuroscience from which he dissents. Here at least you can see what he means, and the history is very interesting. Another problem is that large philosophical themes like dualism and free will are clearly intended to be illuminated by the discussion, but nothing is spelled out. I think Glimcher is mixing up two books. Sometimes it reads as though he's trying to provide a general, user-friendly but rigorous, introduction to neuroeconomics; and sometimes he seems to be advocating his own particular vision of what economics can offer neuroscience.

So, what is the positive program which Glimcher advocates in place of the approach whose history he has investigated? Glimcher assumes that cognition is modular, without ever really spelling out what he means by that, beyond functional specification. He thinks that the modules are each designed to solve an adaptive problem. So the problem naturally arises: how do we analyze the problems the modules evolved to solve? Evolutionary psychology, which also sees the mind as a collection of evolved modules, has a notorious answer to this question, which is to try to understand the life lived by our hunter-gatherer forebears and assume that the modules are solutions to the problems that they faced. Glimcher takes a different tack, arguing that "mathematical theories of decision making that include probability theory must form the core of future approaches to understanding the relationship between behavior and brain" (p. 177). Glimcher takes this ahistorical approach to understanding adaptation because of an overriding commitment to the idea that the nervous system has evolved to make decisions that maximize inclusive fitness. Although he doesn't say so, this is reminiscent of Dennett's intentional stance, but with formal theories of behavior replacing Dennett's intuitive folk psychology as the basis for predictions. (This makes perfect sense if, as theorists like Jon Elster (1989) have argued, we should think of rational choice theory as a mathematically sophisticated development of the underlying structure of folk psychology.) So it's economics, not evolutionary history, that's the means "for defining the problem that an animal faces or the goal that it should achieve" (p. 201).

Now, as Glimcher recognizes, his conception of his project places him squarely in the recent mainstream of behavioral ecology and biomechanics. Quantitative optimality analyses for organisms have been around for years, drawing on engineering to predict things like the optimum structural properties of tubular bones, as well as game theory to predict behaviors such as intraspecies conflict and sex ratios. Glimcher uses examples drawn from foraging theory, where optimality analyses have long been utilized effectively, and Maynard Smith's evolutionary game theory. His picture is one in which the existing formal methods of behavioral ecology are extended into the brain via a search for properties of neurons that can be modeled in economic terms.

Traditional approaches in ecology and population genetics were highly idealized, assuming that phenotypes could be described game-theoretically as strategies that one could then model as genotypes. Everyone knew that this picture was only a simplification—people don't play just one strategy all the time and there is no simple one-one relation between genes and strategies. The picture survived because it made for tractable models and useful predictions. Glimcher, though, wants to use the strategy to understand neural mechanisms. He doesn't look for the genes underlying optimal rational behaviors, but the neurons. And indeed, he argues that he has discovered neurons in parietal area LIP that compute expected utility, and do so in the same basic fashion—corresponding to traditional canons of rational choice—regardless of whether or not the task at hand involves a predictable environment. Glimcher thinks this is a very big deal indeed, because he treats the distinction between predictable and unpredictable behaviors as the distinction between behaviors that, in the Cartesian tradition, can be understood by assuming the organism is an automaton, versus those that can only be understood by positing "the soul or free will" (p. 337).

Glimcher thinks game theory offers a breakthrough here by showing that a purely physical organism can engage in unpredictable behavior. He assumes that hitherto physicalists have argued that a deterministic system must be predictable. But philosophers have been saying for years that a physically determined system can be unpredictable. So I'm afraid Glimcher adds nothing to the philosophical debate. But there are some issues that Glimcher raises at the boundaries of philosophy and the cognitive and social sciences that are of great interest (and his account of his experimental work is fascinating on its own). I'll end by mentioning a couple.

Essentially, Glimcher treats utility, which can be computed by the brain, as a proxy for biological fitness, which cannot. After all, organisms do not strive to maximize fitness directly. Rather, they look for food, mates, predators, etc. If they manage these tasks well enough, they leave more descendants. Glimcher's basic idea is that all this diversity of behavior can be boiled down, at the level of proximate mechanisms, to the operations of brain systems which compute expected utility. This allows utility maximization to be used as a general measure of optimality, instead of specific measures for specific traits. Although many behaviors, especially in other organisms, might be approachable in this way, it's very unlikely that it can provide a general approach to human behavior. It might work when it comes to satisfying preferences (e.g., sex, food) that might have been instilled in us by evolution. But it still seems implausible in light of stresses that have recently afflicted traditional economic models of rationality, that the approach can be extended to more complex human behaviors, where the relevant probabilities may not be ones we are evolved to learn.

There may still be a moral here for the study of complex behavior. Glimcher's orientation on cognition is taken from David Marr. He borrows from Marr the idea that we should ask what the goal of a behavior is, how it is to be modeled computationally, and how it is realized in brain tissue. Yet, the final program does not seem to be Marrian at all. Rather, it relies on defining the problem in terms of expected utility theory and then going straight to the neurons—connecting, as Glimcher puts it, behavior and neurophysiology through a mathematical corpus (p. 319). The middle Marrian level, of cognitive architecture, is not involved. So in so far as the Marrian picture distinguishes the traditional approach to cognitive science, I think it fails to fit the picture in neuroeconomics, or indeed social and cognitive neuroscience more generally. Cognitive psychologists have discovered many important and fascinating effects, but attempts to uncover an independent layer of computational generalizations haven't gotten very far, and indeed there's no consensus when it comes to the computational approach that's needed: classical, dynamic, and connectionist models all have their constituencies.

Perhaps the lesson to take from neuroeconomics is that we can bypass these computational controversies. We have increasing knowledge of the brain, and increasing knowledge of our psychology, including formal economic models that can be rendered more realistic by experimental and observational input from psychology and the other social sciences. Both the social sciences (including parts of psychology) and the neuroscience (including parts of psychology) seem to be more robust than purely Marrian cognitive psychology. In particular, attempts to provide computational models of rational processes and their frailties seem to have gone nowhere. So maybe all we need is the interaction between these two lines of inquiry. Perhaps the attempt to find a general computational approach independent of the specifics of the human brain is just misguided. Rather than trying to understand cognition as a set of abstract computational entities, we should focus on understanding human behavior and linking it to the neurons. Jerry Fodor has long argued that the relations between beliefs and utilities could not be understood using computational tools, and concluded that cognitive science could not, therefore, be done. Fodor's pessimism was based on his belief that computational models of non-demonstrative inference were impossible, but doing things Glimcher's way suggests that such models may not be needed for understanding cognition after all.

Put another way, maybe the three-level approach we got from Marr was over-complicated, and the future of cognitive science will be social science at the top, brain science at the bottom, and nothing in between. This is the project that Glimcher really points us towards, I think, and this book has some fascinating suggestions about how we might start doing it.

References

Elster, J. (1989). *Nuts and bolts for the social sciences*. Cambridge, England: Cambridge University Press.

Dominic Murphy
Division of Humanities and Social Sciences
California Institute of Technology
MC 101-40
Pasadena, CA 91125, USA
Email: murphy@hss.caltech.edu

Early Category and Concept Development: Making Sense of the Blooming, Buzzing Confusion

David H. Rakison & Lisa M. Oakes
New York: Oxford University Press, 2003
442 pages, ISBN: 0195142934 (hbk); $49.50

At the dawn of modern science, philosophers argued about the age-old question of whether humans construct reality from their senses or whether reality is shaped by innate properties of the human mind. *Early Category and Concept Development* adds a new perspective to this question using the lens of developmental psychology.

As Francis Bacon suggested, infants and children might offer an interesting test case for understanding reality through science: They are unspoiled by years of experience, "uncorrupted by false notions" (Cranston, 1967, p. 235). One can thus ask how infants view the world. Do they start with adult-like categories of objects and events? Or do they construct their categories and concepts of objects and events by assembling perceptual regularities witnessed through years of experience? In this superbly edited volume, Rakison and Oakes take us through multiple theoretical perspectives and offer a panoramic view of the kind of data that developmental psychology can add to these questions.

This book continues the philosophical conversation and raises three questions of interest to modern-day scientists. The first concerns constructivism, and is cast in terms of concepts and categories: Do infants have *a priori* conceptual information that guides the ways in which they see and represent their world, or do they use perceptual building blocks to construct object and event categories? The second question concerns the role of language in the development of concepts and categories: Do conceptual categories guide language learning? Alternatively, does language invite or shape the way our categories are formed (Waxman, ch. 9)? The third question—raised by Oakes and Rakison (ch. 1)—asks whether the methods used in developmental psychology are sufficient to address the questions being raised. That is, do the methods employed by psychologists tap into infants' existing and changing knowledge about categories and concepts or do these methods themselves impose a grid that prompts infants to construct a view of the world "online" (Newcombe, Sluzenski, & Huttenlocher, 2005)?

The first eight chapters are devoted to the question of constructivism. That categorization and conceptualization are central to the human experience is taken as a given in these chapters. As Oakes and Rakison note,

> Forming categories reduces demands on our inherently limited memory storage and perceptual processes, and without it we would have to remember independently the same or similar information about each individual member of a category . . . we would have to remember that a particular Rottweiler has four legs and so does a particular poodle. (p. 4)

How do these categories and concepts come to shape our experience? Most of the authors in this volume advocate a "bottom-up" answer to this question, holding that infants attend to perceptual elements in the environment and build concepts and categories through a kind of statistical assessment of the input. Younger (ch. 4), for example, showed infants of different ages sets of stimuli that were systematically varied (e.g., by the size of the body, tail, legs, and ears). She asked whether infants notice experimenter-controlled correlations among these different perceptual elements. The youngest infants in her studies ($\approx$ 3 months) noticed only specific features, while older infants ($\approx$ 9–10 months) detected correlated features that later form categories. Quinn's (ch. 3) research presents a similar developmental story. Infants begin with meager attention to perceptual specifics, only later building complex categories for spatial concepts (e.g., BETWEEN).

These authors support a perceptual, constructivist view. The infants' view of reality differs significantly from the adults' and requires little or no *a priori* organization from an innate human mind. While this is attractive from a parsimonious view, it is not without its critics. First, one can ask whether all percepts are created equal or whether some are more privileged or more noted than others. This traditional *relevance* problem haunts the "bottom-up approach," and is one that Rakison (ch. 7) attempts to address. By way of example, Rakison attempts to solve the relevance problem by proposing that infants prioritize certain percepts over others (i.e., movement of object, object size). There is, then, some organization imposed on perceptual processes used to construct categories and concepts. Second, some authors like Mandler (ch. 5) suggest that infants bring more to the table than just the ability to statistically assemble perceptual information.

Mandler's top-down approach thus stands in stark contrast to that presented by others in this volume. She advocates that instead of moving from specific percepts to broad categorizations, development progresses from broad generalizations to specific instances. She writes:

> . . .infants generalize broadly on the basis of abstract conceptualizations first and only with experience learn to pay attention to perceptual detail . . . even though infants must use various physical features to tell animals such as dogs and cats apart, they do not rely on them when they are construing the meaning of an event and generalizing from it. (p. 115)

The test case for Mandler comes in a design that pits perceptual information against deeper conceptual information. For example, Mandler might show children that a set of toy vehicles (e.g., cars) can be associated with a "key" and that a set of toy animals (e.g., dogs) can be fed with a "bottle." She and her colleagues then ask whether, when presented with a two perceptually similar toys that cross ontological boundaries, infants make the proper inductions or generalize by perceptual similarity. Mandler finds that 14-month-olds will feed a bottle to a bird, but not to a perceptually similar airplane. She concludes that infants therefore go beyond the surface properties of objects to induce something about conceptual *essences*. Gelman's work (ch. 13) presents similar data to suggest that infants are not driven by perception alone, but rather that they operate with a richer set of categories and concepts that allow them to discover deeper order in the world in which they live. For these authors, conceptual biases drive perceptual organization rather than the other way around (cf. Rakison, ch. 7).

The first half of the book offers the most interesting discussion of these foundational issues in psychology and philosophy. It is capped by Cohen's (ch. 8) very balanced commentary on the question of whether there exists a single categorization system (i.e., one that relies on perceptual information) or a dual categorization system (i.e., one that relies on both perceptual and conceptual information). While he favors the bottom-up perspective on these issues, he also admits:

> it is inevitable that unresolved issues will emerge in a field as active and significant as infant categorization. What is most exciting is that those issues are generating empirical questions both within and beyond the traditional realm of infant categorization research. (p. 207)

While the first half of the book asks how infants interpret reality in a nonlinguistic world, the second half (ch. 9–16) explores language's role in the development of categories and concepts. Smith and colleagues (ch. 11) believe that children's early word learning helps create their category knowledge by making some features of objects more salient than other features. For example, children have a bias to extend names for animate objects to objects of similar shape and texture, and names for substances to objects of the same material. Waxman (ch. 9) contends that words serve as an invitation for categorization and heighten attention to common properties shared by a large number of objects. She suggests that infants first attend to category-based commonalities when they hear a word (e.g., animal and vehicle) and only later attend to object properties (e.g., color and texture) and relations. Older infants move beyond objects to link novel adjectives to specific properties of the object they see (e.g., *purple* horse).

The prior papers explore ways in which labels serve as attractors for categories. Other chapters examine the idea that labels provide rich mental structures for infants that tap into underlying theories of conceptual organization (e.g., Gelman & Koenig, ch. 13; Gopnik & Nazzi, ch. 12). According to Gelman and Koenig, "language may be crucial for anchoring kinds" (p. 351). Language, then, does more than merely draw attention to shape or objects. It helps children coalesce multiple sources of information, including function, ontology, intention, and non-obvious properties as they form broader categories and concepts that tap into richer conceptual theories. Linguistic labels become the "placeholders for a theory about why two things are members of the same kind and can be used to make inferences when kind is relevant to the inference" (Markman & Jaswal, ch. 15, p. 400).

In the treatment of both the non-linguistic and linguistic theories of infant development, this volume represents theoretical divisions in the field and offers commentary to place these divisions within the broader purview of psychology. The book also displays a spectrum of different methodologies used in the field of infant development and demonstrates how these methods often generate different answers to questions about the bottom-up and top-down nature of category and concept development. In this sense, the book also serves as a tutorial in burgeoning methods within a field. The reader will be impressed by the ways in which developmentalists peer into the mind of the pre-linguistic child. Readers might also wonder about a hotly debated issue within the field of child psychology: whether the methods used to examine preexisting categories in infants are themselves responsible for the organization that we see in early development (Newcombe et al., 2005).

In the end, then, we can ask whether the study of infants helps to resolve deeper philosophical questions about how humans experience reality. The developmental literature does offer some interesting data on this question. Minimally, the chapters in this volume show that infants are able to find constancy in their world and to detect correlations of perceptual stimuli. The data also demonstrate that infants can form categories of both objects and events from these more meager perceptual beginnings. A number of researchers hold that this alone suggests we can now offer a parsimonious path from perception to conception.

Do infants have more than perceptual atoms to work with? The evidence that babies seem to be making inductive inferences and generalizations about the objects and events in their world is interesting and tantalizing. Whether these data are convincing, the reader must evaluate. Wherever one stands in the theoretical debate, this clearly written and well-edited book adds to the chorus of opinion that might inform answers to those questions plaguing the study of reality and its construction. This book is a must-read for those who study human categorization and conceptualization. It offers an up-to-date and well-argued compendium of opinion and research from the finest minds in developmental psychology, and illuminates the current status of the empiricist-rationalist debate within psychology.

References

Cranston, M. (1967). Francis Bacon. In P. Edwards (Ed.), *Encyclopedia philosophy* (Vol. 1, pp. 235–240). New York: Macmillan.

Newcombe, N. S., Sluzenski, J., & Huttenlocher, J. (2005). Preexisting knowledge versus on-line learning. What do young infants really know about spatial location? *Psychological Science*, *16*, 222–227.

Shannon M. Pruden, Kathy Hirsh-Pasek, & Julia Parish
Department of Psychology
Temple University
1701 North 13th Street
Philadelphia, PA 19122-6085, USA
Email: spruden@temple.edu