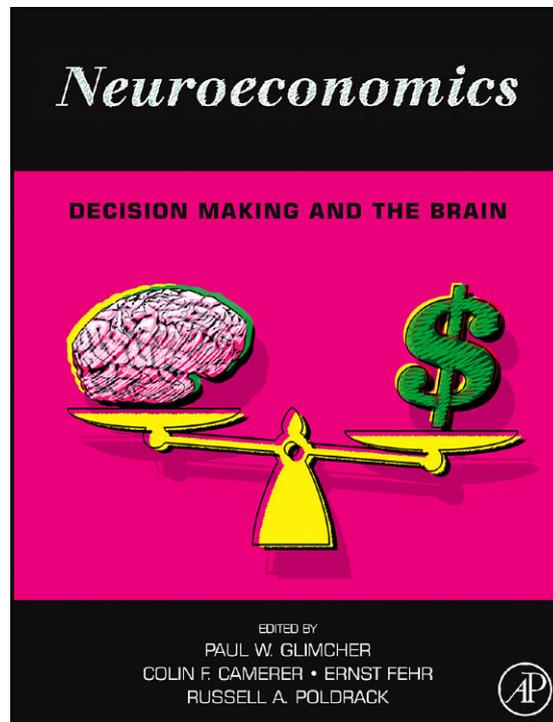


**Provided for non-commercial research and educational use only.
Not for reproduction, distribution or commercial use.**

This chapter was originally published in the book *Neuroeconomics: Decision Making and the Brain*, published by Elsevier, and the attached copy is provided by Elsevier for the author's benefit and for the benefit of the author's institution, for non-commercial research and educational use including without limitation use in instruction at your institution, sending it to specific colleagues who you know, and providing a copy to your institution's administrator.



All other uses, reproduction and distribution, including without limitation commercial reprints, selling or licensing copies or access, or posting on open internet sites, your personal or institution's website or repository, are prohibited. For exceptions, permission may be sought for such use through Elsevier's permissions site at:

<http://www.elsevier.com/locate/permissions>

From: *Neuroeconomics: Decision Making and the Brain*
Edited by Paul W. Glimcher, Colin F. Camerer, Ernst Fehr and Russell A. Poldrack
ISBN: 978-0-12-374176-9

© Copyright 2008 Elsevier Inc.

Academic Press.

Author's personal copy

P A R T V

THE NEURAL MECHANISMS
FOR CHOICE

The Neural Mechanisms that Underlie Decision Making

Charles R. Gallistel

OUTLINE

| | | | |
|---|-----|------------|-----|
| Introduction | 419 | Matching | 422 |
| Can/should Neurobiological Data Constrain Economic Theorizing? | 420 | References | 424 |

INTRODUCTION

The first word that springs to mind in introducing contemporary work on the neurobiology of decision making is *ignoramus*: we do not know the neural basis of economically consequential decision making. However, *ignoramus* does not imply *ignorabimus* – that we *will* not know. Not knowing *now* is not a reason to despair of knowing, but rather the source of our aspiration to know. The chapters in this section are an introduction to contemporary efforts to turn that aspiration into knowledge.

The task is daunting, because it involves the establishment of empirically well-secured hypotheses linking phenomena and entities at different levels of analysis. Three levels of analysis come into play in

neuroeconomics: the economic, the psychological, and the neurobiological. The objects of analysis and the terms in which the analyses are conducted at one level have no obvious referents or definitions at the other levels. The economist works at the market level, trying to understand how the decisions made by individuals in a market economy determine quantities such as aggregate supply, aggregate demand, and price. These are definable only by reference to a market economy. They are not the variables that psychologists work with in trying to understand the behavior of human and non-human animals, because they do not operate *within* individual subjects. Neither are they attributes of the behavior of an individual subject. On the other hand, economists do posit variables (for example, subjective utilities) and processes (for example, decision processes) that they suppose do operate

within the brains of individual subjects to determine the decisions that collectively create demand, help to determine supply, and establish price. However, these psychological variables and processes have, in their turn, no obvious referents at the neurobiological level. Indeed, in some neurobiological circles there is doubt whether symbolic variables such as subjective values have a role to play in understanding how the brain works (Edelman, 1989).

An illustration of the difficulty of establishing, to general satisfaction, hypotheses that link concepts defined at one level of analysis to concepts defined at a different level is provided by the lengthy history of attempts to link the concept of a nerve impulse to the concept of an action potential. The concept of a nerve impulse is defined by reference to behavioral observations. It has been invoked, at least since the time of Descartes, to explain the fact that pinching the toe gives rise to contractions in muscles far removed from the toe. The concept of an action potential is defined by reference to electrophysiological observations on nerves and muscles, which were first made by du Bois-Reymond in 1848 (see du Bois-Reymond, 1852). He believed that the brief alterations in the electrical potential between the insides of (a bundle of) axons and their outsides that accompanied the behavioral action of a nerve was the physical realization of the nerve impulse. The physical nature of the impulse had heretofore been an object of much speculation, with the vitalists (then still an influential school of thought) doubting that it had a physical realization. Du Bois-Reymond was, of course, right; no one doubts this linkage hypothesis now. However, in 1936, 40 years after Du Bois-Reymond died, the Nobelist A.V. Hill opened a review of experimental work on excitation and accommodation in nerves with the assertion that the physical identity of the nerve impulse remained an unsolved problem (Hill, 1936).

I stress the difficulty of firmly establishing these linkage hypotheses in order to dash one hope that economists might entertain about the fruits of neuroeconomic research: measures of neural signaling activity are unlikely to place hard-to-measure economic decision variables like subjective values on a more secure foundation in our lifetimes. In order for that to happen, we would need to establish a secure hypothesis linking subjective values to measurable neurobiological quantities. We would have to believe, for example, with a high degree of confidence, that subjective values are physically realized by the firing rates of identifiable neurons and that therefore, as Glimcher boldly hypothesizes, "they take as their natural units action potentials per second" (see Chapter 32, page 509)

CAN/SHOULD NEUROBIOLOGICAL DATA CONSTRAIN ECONOMIC THEORIZING?

Closely related to my skepticism regarding the prospects for replacing messy behavioral measures of critical subjective variables with precise neurobiological measures are my doubts about the ability of neurobiological data to constrain economic theorizing. The assumption that neurobiological data can and should constrain both psychological and economic theorizing is widespread and seldom questioned. It is, for example, a foundation for neural net theorizing of all kinds. In Chapter 28, Rangel argues that "A central goal of neuroeconomics is to construct theories of value-based decision making that are neurobiologically grounded," and that, "As in the case of perceptual decision making, a key test of the model is the presence of neurobiological signals that resemble the computations that it makes." And Glimcher argues in Chapter 32 that "The explicit ties to neurobiological data will reveal that only a tiny space in the vast landscape of economic theory can be viewed as compatible with the human neuroarchitecture."

In order for neurobiological data to constrain psychological and economic theorizing, one must be persuaded beyond reasonable doubt that the neurobiological variables that were measured or the neural structures that were described were actually correctly linked to a psychological variable, such as subjective value, in a valid theory of the psychological causation of the behavior. If the neurons whose firing rates are measured are not the neurons whose firing rates physically realize the subjective values in question, then the measures made on them are irrelevant to the psychological or economic theories in which those variables play a role. More radically, if, at the point in the brain where decisions are made, subjective values are not represented by firing rates at all but rather by some other measurable neurobiological variable – for example, the concentrations of intracellular signaling molecules – then any measure of firing rate or of quantities derivative of it, such as BOLD signals, are irrelevant to our understanding of the decision-making process. Moreover, if economically consequential decisions are not made in the cortex but elsewhere – for example, in the hypothalamus – then the architecture of some portion of the cortex is irrelevant. The architecture of the hypothalamus is radically different from the architecture of the cortex, and neuroscientists are more than a little unclear about the functional relevance of the architecture of either one. There is little doubt that their architecture is an important part

of their functional capability – why should the brain be different from the liver, or the eye? – but just what their capabilities are and just how the architecture supports them are far from decided questions. In many cases, there are not even any clear hypotheses.

In short, the relevance of neurobiological data to psychological and economic theorizing depends on the strength of the hypotheses that assert that specified neural variables are the physical realization of specified psychological and economic variables. These hypotheses are among the most difficult to establish in all of science, and there are very few truly secure ones in contemporary behavioral neuroscience. Locutions such as “might be involved in,” which recur often in behavioral neuroscience, are an indication of the weakness of most of its linkage hypotheses.

Moreover, and perhaps most tellingly, I believe that any study of successful physical reductionism – that is, any study of the establishment of hypotheses such as that the action potential is the physical realization of the nerve impulse, or that base-pair sequences in DNA are the physical realizations of genes – will show that the reductionist program only succeeds when the variable whose physical identity is sought is already well defined and understood in its own terms, at its own level of analysis, within a theoretical framework that rests on observations at that level. As A.V. Hill's (1936) review shows, neuroscientists acquired an elaborate qualitative and quantitative understanding of the nerve impulse based on observation of its behavioral effects (muscle twitches), without recourse to non-behavioral observations. Similarly, we had an elaborate qualitative and quantitative knowledge of genes, based on the study of the inheritance of phenotypes, prior to their identification with base-pair sequences. Indeed, it is still the case, and I would argue it always will be, that the natural unit for measuring the distance between two genes is the centimorgan, not the nucleotide.

A corollary of my beliefs about the preconditions for successful physical reductionism is that neuroscientists have more to gain from economists and psychologists than the latter have to gain from the former. I see neuroeconomics as the attempt to discover the neural variables and structures that well-established economic and psychological theory tell us should be there. Insofar as economic and psychological theory do not rest on secure foundations sunk in the soil of observations native to those sciences, then behavioral neuroscientists must join in the effort to buttress them through behavioral experiments, of which this section contains many excellent examples.

Linkage hypotheses are also called bridge hypotheses, because they bridge between levels of analysis. A

bridge is never stronger than its weakest buttress. If at its psychological end one finds a conceptual swamp, then there is not much prospect for strengthening it by neurobiological observation, because the weaker the bridge, the greater the question as to the relevance of those observations. It is in this spirit – to what extent can we find in the nervous system the variables and processes that economic and psychological theory say must be there? – that I believe the chapters in this section are most profitably read. An example of the influence that well-founded psychological theory has had on neuroeconomic research is the influence of the drift diffusion model of decision making (see Ratcliff and Smith, 2004) evident in many of the chapters in this section.

Neuroeconomics rests strongly on the assumption that the variables and decision-making processes that are economically consequential are found in non-human animals as well as in humans. I think this assumption will prove to be true, but it is only an assumption, and there are some observations that might give one pause before embracing it. In particular, there is little reason to think that anything resembling a market economy exists in any non-human context. A corollary is that it is unclear that prices, whose natural unit is money – a completely fungible commodity – play any role in the behavior and mental processes of any non-human animal.

Money greatly facilitates human exchange, because two different goods are often not interchangeable in many contexts. A king in need of a horse may prefer a broken-down nag to all the tea in China. Shizgal and his collaborators (for review, see Shizgal, 1997) point out that a lack of interchangeability is also a critical issue in an animal's individual economy. Animals need many different things – and the things they need are not interchangeable. No amount of sexual titillation will substitute for orange juice in the maintenance of a monkey's internal milieu, and no amount of orange juice will substitute for an encounter with a receptive female in furthering a male monkey's reproductive success. So how is it possible to get a monkey to give up orange juice in order to be allowed to look at erotic pictures, as Platt and his collaborators manage to do in the experiments described in Chapter 29?

Shizgal and his collaborators trace the origin of mechanisms that compute subjective utility to this problem inherent in the animal economy: generally speaking, an animal can only pursue one goal at a time. The different goals it might pursue are incommensurable. In order to decide which of several incommensurable goals to pursue, the factors favoring one goal or another must be reduced to a common evaluation on a unidimensional scale, the scale of

subjective utility. Shizgal and his collaborators suggest that in the course of computing the subjective utility of positive behavioral outcomes, the brain computes an intensity of reinforcement signal. They identify the intensity of a reinforcement signal with the signal that mediates the rewarding effects of focal electrical stimulation of the medial forebrain bundle. This signal is carried by myelinated axons in the medial forebrain bundle, which runs through the lateral hypothalamus, connecting the ventral midbrain to the frontal lobes (Shizgal and Murray, 1989).

The intensity of reinforcement is not the same as the utility of an outcome, because the utility is affected by the rate at which reward can be obtained and the delay in obtaining it, among other things. Nonetheless, Shizgal and Murray argue that the intensity of reinforcement is like subjective utility in that it is abstracted from the sensory/perceptual qualities inherent in actual rewards, like orange juice and erotic pictures. Intensity of reinforcement is a dimension along which sexual pictures vary as one moves from those seen in textbooks on sexual hygiene to those seen in the centerfold of a men's magazine, and along which orange juice varies as one goes from an overly diluted restoration from frozen concentrate to freshly squeezed. They argue that there is a pure intensity of reinforcement signal that represents variation along this intensity dimension independent of the other attributes of an outcome, and identify this intensity of reinforcement signal with the rewarding signal in brain stimulation reward. They argue that the electrodes in the medial forebrain bundle stimulate the system that computes the utility of outcomes after the point at which the current value of those outcomes has been computed. Moreover, they show that conventional outcomes, like saline and sucrose (non-interchangeable solutions) sum their reinforcing intensities with the signal produced by the brain stimulation reward, and that reducing or eliminating hunger and salt-deprivation reduces or even reverses the intensity of reinforcement from sucrose and saline, but neither manipulation reduces the intensity of reinforcement of the brain stimulation reward.

Their hypothesis about the natural function of the neural substrate for brain stimulation reward explains one of the most salient characteristics of the brain stimulation reward phenomenon: the lack of satiation. For any natural reinforcement, enough is enough, because the reinforcement reduces the conditions that gave it value; however, for brain stimulation reward, enough is never enough. No matter how much brain stimulation reward the animal has already obtained, its appetite for more is undiminished, as would be the case for an outcome that had a high positive intensity

of reinforcement no matter what the animal's physiological state.

I mention their Shizgal hypothesis in part because it is one of the more fully studied and experimentally supported hypotheses in the field of neuroeconomics. It is, I believe, the only one that rests on a demonstration that a signal artificially injected into the brain is causally effective in a manner consistent with the function ascribed to the stimulated system. The demonstration of the causal efficacy of the signals identified with a causal variable in behavior is a *sine qua non* for a secure neurobehavioral linkage hypothesis.

Another reason for describing this hypothesis and the evidence for it is that it takes us out of the cortex, where most neuroeconomic research focuses, down deep in the brain to the hypothalamus. Curiously, the dopamine neurons that carry the timed reward-expectation signals described in several chapters in this volume (see also Schultz *et al.*, 1992; Fiorillo *et al.*, 2003) also run in the medial forebrain bundle. More curiously still, drugs that block the D2 receptor, a postsynaptic receptor molecule in dopamine signaling pathways, block the reinforcing effect of brain stimulation reward, in relative doses predicted by their relative affinity for that receptor (Gallistel and Davis, 1983). Among the many things on the agenda of neuroeconomic researchers in coming decades must be understanding how these hypothalamic mechanisms relate to the cortical mechanisms studied by electrophysiological means in many of the chapters in this and other sections.

MATCHING

I close this introductory essay by returning, in connection with some recent work from my own laboratory, to my point about the importance of having secure psychological theories as a precondition for the neurobiological identification of the variables in those theories and the mechanisms that realize them. Among the requirements for a valid psychological theory of choice in a given paradigm is that it correctly identifies the choice as the subject represents it. A second requirement is that it correctly describes the inputs to the decision variable, and the manner in which those inputs are processed so as to yield a decision. In Chapter 31, Lee and Wang apply their neurobiologically based stochastic model of decision making to, among other things, the matching phenomenon, which is also touched on in Chapters 30 and 32, from the Newsome and Glimcher labs, respectively, both of which have studied this phenomenon

from a neuroeconomic perspective. An assumption in most models of this phenomenon is that the animal's nervous system represents this as a choice between two possible options. Other assumptions are that (1) its choice is based on a running average of (2) the outcomes of its recent choices. There are reasons to doubt all three of these assumptions.

To begin with, most models of matching in the neuroeconomic literature assume discrete trials, because that is how the matching paradigm is run in this literature (for good technical reasons). However, life is not organized into trials. I believe that the real test of a model of a matching model is, first, whether it can even be applied to the free-operant matching paradigm (Herrnstein, 1961), which, as its name suggests, does not impose a trial structure on the animal's experience and responses. In this case, there are two places to forage, and either of them occasionally yields some food; the subject is free to apportion its time between the two locations however it sees fit. Subjects move back and forth at will between the two locations, apportioning their time such that the ratio of the expected durations of their stays at the two locations matches the ratio of the expected incomes from the two locations (food per unit time). Many trial-based models fail this first test, because it is unclear how to generalize them to the continuous time (no trials) situation, which must be regarded as the more natural situation.

Secondly, the paradoxical results from experiments in which subjects are first observed to match given each of two different pairs of locations (with income ratios of 1:2 and 2:4), and are then offered a choice in probe trials between a pair composed of one option from each of the two original pairs (e.g., the options with incomes in the ratio 2:2), imply that the subjects are not in fact choosing between the two options on offer. They are not making the choice they appear to be making. Rather, they are deciding when to abandon whichever option they have already chosen (Belke, 1992; Gibbon, 1995; Gallistel and Gibbon, 2000; Gallistel *et al.*, 2001).

Thirdly, and of particular interest to economists, the results from our experiments are not consistent with the assumption that matching is based on the gains from these choices – that is, on the relative profits from the behavioral investments. Critical here is the distinction between *income* and *profit* – a distinction very familiar to economists, but less so to psychologists. In free-operant matching, the income from a location is simply the amount of food obtained there per unit of time (full stop; note, *not* per unit of time spent there), while the profit is the income per unit of time spent there. Matching equates profits,

not incomes. Traditionally, models have assumed that matching was the slow equilibration of profit by a hill-climbing process in which the investment ratio was adjusted until the profit ratio was 1:1 (Herrnstein and Vaughan, 1980). We have shown, however, that the process of adjustment can be extremely fast; it can go to completion within the span of a single inter-reinforcement interval at one of the locations (Gallistel *et al.*, 2001). Moreover, when computed on a visit-by-visit basis, profit is a much noisier variable (because income appears in its denominator, and its numerator, visit duration, is an exponentially distributed random variable). Thus, changes in income are apparent much sooner than are changes in profit. We showed that the adjustment to a new ratio of incomes often occurred before there was any evidence of a change in the relative profitability of visits to the two locations (Gallistel *et al.*, 2001).

The abruptness of the adjustment to changes in incomes is a challenge to models that assume a decision variable that is a running average of recent outcomes, whether those outcomes are construed as profits or incomes. Running averages cannot change from one extreme to the other between one datum and the next. An alternative model of the decision process is one in which the estimates of the incomes are based on small samples, with re-sampling only when an analysis of the datastream that looks for changes in income finds a change (Gallistel and Gibbon, 2000; Gallistel *et al.*, 2007). With such a procedure, the successive samples contain no overlapping data, so the estimate of an income ratio can change by any amount between one datum and the next.

The assumption that matching is based on income, not profit, is an extremely odd one, because it assumes that the animal does not take into account the impact of its own sampling behavior on the incomes that it observes. If it very rarely visits a high-profit location, then its income from that location will be low – lower, perhaps, than its income from a frequently visited low-profit location. This oddity has predictable consequences, which are even odder. It predicts that when the incomes from two locations are made strictly dependent on the subject's investments in them – i.e., on how long it spends at each location – then there should be an almost exclusive preference for the more profitably location (so far, so good), BUT it should be unstable: every now and then subjects should abandon the more profitable location for a prolonged stay at the less profitable location, even though in doing so they immediately reduce both their overall income and their overall profit. That prediction should make any economist sit up and take notice. I dare say most economists regard it as so unlikely that it is not worth

running the experiment. We did run the experiment, and that is what we in fact observed (Gallistel *et al.*, 2007). Moreover, these paradoxical reversals in preference were again extremely abrupt; they were not consistent with the assumption that preference was based on a running average of recent outcomes.

If the study of matching and its neurobiological basis is to be an important part of neuroeconomic research – and I believe it should be – then we have got to get the psychological model right. If we go looking for a running average of the profits, and if matching is not based on such a quantity, then we will look in vain. Worse yet, we may find it, and we will then erroneously assume that finding such a signal in the nervous system validates our psychological model.

References

- Belke, T.W. (1992). Stimulus preference and the transitivity of preference. *Animal Learn. Behav.* 20, 401–406.
- du Bois-Reymond, E. (1852). *Animal Electricity*. London: Churchill.
- Edelman, G.M. (1989). *Neural Darwinism*. Oxford: Oxford University Press.
- Fiorillo, C.D., Tobler, C.D., and Schultz, W. (2003). Discrete coding of reward probability and uncertainty by dopamine neurons. *Science* 299, 1898–1902.
- Gallistel, C.R. and Davis, A.J. (1983). Affinity for the dopamine D₂ receptor predicts neuroleptic potency in blocking the reinforcing effect of MFB stimulation. *Pharmacol. Biochem. Behav.* 19, 867–872.
- Gallistel, C.R. and Gibbon, J. (2000). Time, rate, and conditioning. *Psychological Rev.* 107, 289–344.
- Gallistel, C.R., Mark, T.A., King, A.P., and Latham, P.E. (2001). The rat approximates an ideal detector of changes in rates of reward: implications for the law of effect. *J. Exp. Psychol. Animal Behav. Proc.* 27, 354–372.
- Gallistel, C.R., King, A.P., Gottlieb, D. *et al.* (2007). Is matching innate? *J. Exp. Anal. Behav.* 87, 161–199.
- Gibbon, J. (1995). Dynamics of time matching: arousal makes better seem worse. *Psychonomic Bull. Rev.* 2, 208–215.
- Herrnstein, R.J. (1961). Relative and absolute strength of response as a function of frequency of reinforcement. *J. Exp. Anal. Behav.* 4, 267–272.
- Herrnstein, R.J. and Vaughan, W.J. (1980). Melioration and behavioral allocation. In: J.E.R. Staddon (ed.), *Limits to Action: The Allocation of Individual Behavior*. New York, NY: Academic Press, pp. 143–176.
- Hill, A.V. (1936). Excitation and accommodation in nerve. *Proc. R. Soc. Lond. B* 119, 305–355.
- Ratcliff, R. and Smith, P.L. (2004). A comparison of sequential sampling models for two-choice reaction time. *Psychological Rev.* 111, 333–367.
- Schultz, W., Apicella, P., Scarnati, E., and Ljungberg, T. (1992). Neuronal activity in monkey ventral striatum related to the expectation of reward. *J. Neurosci.* 12, 4595–4610.
- Shizgal, P. (1997). Neural basis of utility estimation. *Curr. Opin. Neurobiol.* 7, 198–208.
- Shizgal, P. and Murray, B. (1989). Neuronal basis of intracranial self-stimulation. In: J.M. Liebman and S.J. Cooper (eds), *The Neuropharmacological Basis of Reward*. Oxford: Clarendon Press, pp. 106–163.