American Economic Association

Is There a Method of Neuroeconomics?
Author(s): Aldo Rustichini
Reviewed work(s):
Source: *American Economic Journal: Microeconomics*, Vol. 1, No. 2 (August 2009), pp. 48-59
Published by: American Economic Association
Stable URL: http://www.jstor.org/stable/25760360
Accessed: 04/01/2013 15:03

Your use of the JSTOR archive indicates your acceptance of the Terms & Conditions of Use, available at
http://www.jstor.org/page/info/about/policies/terms.jsp

JSTOR is a not-for-profit service that helps scholars, researchers, and students discover, use, and build upon a wide range of content in a trusted digital archive. We use information technology and tools to increase productivity and facilitate new forms of scholarship. For more information about JSTOR, please contact support@jstor.org.

American Economic Association is collaborating with JSTOR to digitize, preserve and extend access to *American Economic Journal: Microeconomics.*

http://www.jstor.org

This content downloaded on Fri, 4 Jan 2013 15:03:54 PM
All use subject to JSTOR Terms and Conditions

American Economic Journal: Microeconomics 2009, 1:2, 48–59
http://www.aeaweb.org/articles.php?doi=10.1257/mic.1.2.48

Is There a Method of Neuroeconomics?†

By ALDO RUSTICHINI*

This note tries to state, precisely, the method of neuroecomics, and is based on the discussion in B. Douglas Bernheim's (2009) *appraisal. We claim that the theory formulates hypotheses modeling the choice process as an algorithmic procedure. The hypothesis of the algorithmic procedure imposes restriction on the neural processes implementing it, and, so, a joint test of the hypothesis based on behavioral and neural data is possible, increasing the statistical and the explanatory power of the theory.* (*JEL* B41, D87)

I. Methodological Frameworks

The discussion of B. Douglas Bernheim (2009) introduces two different methodological frameworks.

Neural Evidence and Economic Predictions.—The first framework is familiar. There is a variable y to be explained or predicted, which is the economic choice or behavior. There are variables, like income or taxes, denoted by $\mathbf{x}$, that are potential explanatory variables. Economic analysis is an attempt to identify a probability $\eta(\cdot \mid x)$ on the space of observed variables, conditional on the observed explanatory variables. Neuroeconomics can postulate, in addition to the vector $\mathbf{x}$, an additional vector of unobservable environmental variables $\boldsymbol{\omega}$, working their way to affect y through some variable representing, say, neural activity z, which may be observable by the neuroeconomist.

The effect of the unobserved variables $\boldsymbol{\omega}$ is treated as noise in standard economic analysis, whereas the introduction of the additional observed variable z gives some information on the distribution of $\boldsymbol{\omega}$, and on the form of the function producing the neural activity. Since neuroeconomics adds observable variables to the analysis we have nothing to lose by extending the range of variables we consider, but we may also gain very little. However, if the variables in $\boldsymbol{\omega}$ are sufficiently uncorrelated with $\mathbf{x}$, then the observation and use of the variables z can increase the predictive power of the model.

This may be taken as a weak but reasonable defense of the use of neuroeconomics in economic analysis. I think neuroeconomics has much more to offer.

* Department of Economics, University of Minnesota, Hanson Hall, 1925 4th Street South, 4-101, Minneapolis, MN 55455 (e-mail: arust@econ.umn.edu).

† To comment on this article in the online discussion forum, or to view additional materials, visit the articles page at http://www.aeaweb.org/articles.php?doi=10.1257/mic.1.2.48.

This content downloaded on Fri, 4 Jan 2013 15:03:54 PM
All use subject to JSTOR Terms and Conditions

Algorithms and Neural Structures.—The second methodological framework is more ambitious, and is closer to what neuroeconomists (or at least I) have in mind. It is the method of a research program that tries to give a "mechanistic, behavioral, and mathematical explanation of choice that transcends the explanations available to neuroscientists, psychologists, and economists working alone." (Paul W. Glimcher and Rustichini 2004).

The method is the following. A set of choice correspondences describes all possible economic behaviors. Each choice correspondence c maps feasible sets into choices (elements of the feasible set), and economists try to determine which restrictions (or axioms) the c of an individual obeys. Standard decision theory might indicate a utility function u with the property that, when maximized over the feasible set, F produces the choice $c(F)$. This u is a purely conceptual device, and testing whether the decision maker really selects his choice by maximizing u is a misunderstanding of the method of economic theory.

Differently from "mainstream economics," one can hypothesize that the correspondence c is generated by a computational decision algorithm, a. In turn, this algorithm is implemented by a neural structure, n. Differently from standard functional representations (like the u function) axiomatized in decision theory, the neural structure n is now part of the hypothesis, hence evidence concerning the neural architecture is admissible in hypothesis testing.

Joint Testing.—What is hypothesis testing in this framework? Bernheim (2009) suggests that the hypothesis we test is that choices satisfy an axiom, call it Ax; formally, the hypothesis is that the choice correspondence c belongs to a set C_{Ax} of correspondences that satisfies some axiom Ax. In mainstream economics, the hypothesis is rejected if the observed choices are not compatible with the axiom. In the extended framework suggested by Bernheim (2009), the test uses both neural evidence and behavioral evidence. The observed choices define a set of neural structures that are compatible (that is, that can implement), the choice correspondences that satisfy the axiom Ax, call them N_{Ax}. Evidence from neuroscience selects a set E of possible neural structures. If there are neural structures that are compatible with the behavioral evidence (that is, are in N_{Ax}) and with the neural evidence (are in E), then the hypothesis (that c belongs to a set C_{Ax}) is not rejected.

A. *The Method of Neuroeconomics*

Slightly reformulating the formulation of Bernheim, we can summarize the basic points of the method:

- Choice is assumed to be produced by an algorithm, that is a procedure describing, in detail, every step going from the presentation of the set of options to the subject, up to the selection of one of the options. Each step of the algorithm is simple enough that it is possible to provide a description of its neural implementation. For example, maximizing a function over a set is not a simple step. The algorithm is general enough that it can produce an output for every economic choice problem.

This content downloaded on Fri, 4 Jan 2013 15:03:54 PM
All use subject to JSTOR Terms and Conditions

- The hypothesis we test is that the algorithm has a specific form; that is, the unknown state of the world for the scientist is the algorithm.

- As always in statistical testing, the scientist has experiments available that are maps from the unknown state of the world (the set of algorithms) to signals; that is the experimental evidence. Within a neuroeconomic investigation, he has two types of evidence (signals) available: (i) the choice behavior (which might include the option selected out of each feasible set but also the response time, the error rate, or even speech acts); and (ii) the evidence on the neural structure that is implementing the choices. This includes brain activation in *fMRI* imaging experiments, the effect of pharmacological or *TMS* treatment, and so on.

- The additional statistical power of the experiment derives from the joint restriction on observed variables coming from the two experiments. The information gained may be substantial, and might not be reached with a large number of repetitions of one of the two experiments. For instance, if the distribution on the space of choice signals, given algorithms a^1 and a^2, is nontrivial but small, whereas the corresponding distributions on the space of neural observations are different, then a single neural observation is more informative than many behavioral ones, even if we only care about prediction on choice behavior.

- The additional explanatory and predictive power derives from the specific algorithmic form given to the choice process.

Within the conceptual structure presented here, one way of presenting the objection to the entire research program of neuroeconomics made by Faruk Gul and Wolfgang Pesendorfer (2005) may be the following. If the two algorithms a^1 and a^2 differ very little in the distribution of signals on the space of choice signals, then we should have correspondingly very little interest in the difference, at least as economists. I disagree with this position. The separation between a^1 and a^2 may be useful in prediction out of sample, or in different environments, like the multi-period extension of the choice problem of the observed choice.

How does this method differ form the standard one? The main difference is that the functional representation of choice in decision theory is not considered a testable hypothesis, whereas the algorithmic specification is. This difference gives a new role to the axiomatic method, and does not make it useless. The axiomatic method is a condition for clear thinking. I would like to note that many of the fundamental paradoxes in the "standard theory" (for instance, M. Allais' and Daniel Ellsberg's) have been historically found after, and because a precise theory had been formulated.

I will now argue that the methodological approach suggested here has been in use in neuroeconomics, and has produced some of the most interesting results in the field. I will use one specific example to illustrate the method, and I will try to explain why the method does not only produce a plausible procedure, but also how this procedure can give a disciplined (non-ad hoc) explanation of observed inconsistencies in the decision process.

This content downloaded on Fri, 4 Jan 2013 15:03:54 PM
All use subject to JSTOR Terms and Conditions

II. Inter-temporal Choice

I begin with the workhorse of modern standard analysis of dynamic economies. The theory of decision in an inter-temporal problem has been largely based on the dynamic programming solution of the problem based on the value function associated with the problem. Let S be the set of states, A the set of actions (both finite, for simplicity), and for every pair (s, a), let $T_{(s,a)} \in \Delta(S)$, the probability of the next state. The decision maker receives, for every pair of states and action, a current reward given by $r : S \times A \to R$ payoff function, and discounts the future by a factor δ.

A Familiar "As If" Model.—A rational decision maker that knows the reward function r and the transition function T can solve the problem with standard dynamic programming techniques that determine the value function

$$V(s) = \max_{a \in A} \left(r(s,a) + \delta \sum_{s' \in S} T(s,a)\,(s')\,V(s') \right). \tag{1}$$

The optimal choice of action is defined by the maximization problem in (1). The prediction of the theory is that the decision maker will choose, in every period, such action. This is, of course, an "as if" theory. It does not assume that the optimal action is chosen with a conscious and explicit process according to the maximization problem in (1). It does not explain or assume, in any way, how the decision maker knows the function V. The problem of how the decision maker gets to know this function is not addressed.

The model has proved extremely useful in analyzing several issues in microeconomics as well as macroeconomics. It is widely considered an unrealistic abstraction, however. From the pioneering analysis of Herbert A. Simon (1957), economists have attempted to deal with the lack of realism of the model. The task has proved to be difficult. Several models have been suggested in the last decades, in the research program of bounded rationality. The specific contribution of neuroeconomics is to provide a model that has stronger support from the existing evidence, on the basis of the method outlined here. Let us begin with the specification of the algorithm.

A Computational Decision Algorithm.—An algorithmic description of a process that in the limit produces the value function V and the optimal policy is the following. The decision maker stores an approximation of the value function, and updates it in every k iteration. The function is modified only at the state that occurs in the current period. The values for the other states are left unchanged in that iteration. The adjustment takes the following form. The current state s is observed, the action a is taken (how this action is decided is going to be discussed in a moment), and then the current reward r and the next state s' are observed. The next state s' is chosen by nature according to the probability described by the function T. Note that the observer does not need to know the reward function r or the transition function T. He just needs to be able to observe the current state, the chosen action, the realization of the reward, and the next state. So even if the realized current reward and the next

This content downloaded on Fri, 4 Jan 2013 15:03:54 PM
All use subject to JSTOR Terms and Conditions

state are determined by the two functions r and T, they do not appear in the updating functions below, which only use the *realized* state and reward.

On the basis of the stored value function V_k, the difference between the realized and the expected value can be computed. We denote it by PE, for prediction error, and we define it formally as:

$$PE = r + \delta V_k(s') - V_k(s). \tag{2}$$

Given this "error," the next value function can be determined as:

$$V_{k+1}(s) = V_k(s) + \lambda_k PE. \tag{3}$$

The factor λ_k determines the size of the adjustment in every period. Intuitively, the adjustment is at least qualitatively correct. For example, if the prediction error is positive, then the k^{th} estimate of the value at s was too low, and it needs to be adjusted upward. We can now return to the issue of how the action a is determined. Define the Q_k value of the pair (s,a) (intuitively, the value of being at state s and choosing a, and evaluating the continuation value from the next state according to the approximation function V_k) as

$$Q_{k+1}(s,a) = Q_k(s,a) + \lambda_k[r(s,a) + \delta V_k(s') - Q_k(s,a)]. \tag{4}$$

So, the value of the pair (s,a) is adjusted as the function V_k depending on the current prediction error. Just as in the case of the function V_k, the function Q_k is modified only at the state and action that occur in the current period, while the other values are left unchanged in that iteration.

We can now set the probability of choosing the action a as

$$P_k^s(a) = \frac{\exp(\gamma Q_k(s,a))}{\sum_{b \in A} \exp(\gamma Q_k(s,b))}. \tag{5}$$

The function Q_k, produced at the iteration k, is stored just like the function V_k, and is updated when the information on the next state s' is available. The quantal response function 5 matches the choice behavior well (see, for example, Camillo Padoa-Schioppa and John A. Assad 2006).

The equations (2–5), together with our basic model summarized by the functions r and T, completely describe a stochastic process on value functions, states, and actions. Together they provide a model of adaptive learning that converges to the optimal value function 1 used by economists. Just as the model described by the optimal value function 1, this is an abstract model, but it provides an algorithm that makes explicit how the function V may be learned. The algorithm is known as the actor-critic model (see Richard S. Sutton and Andrew G. Barto, 1998). The *critic* learns to predict the reward associated with a state, and the *actor* keeps information about the rewarding outcomes of actions to guide future choices.

This content downloaded on Fri, 4 Jan 2013 15:03:54 PM
All use subject to JSTOR Terms and Conditions

Compatible Neural Structures.—Is there evidence from research in neuroscience and neuroeconomics to support the view that this is the algorithm followed in making decisions? This evidence has been cumulating for years, and we review it here briefly.

Behavioral evidence supporting the idea that learning follows such a process is provided by the *blocking paradigm*. If a stimulus A fully predicts a reward r, and a stimulus B is then presented together with A, then the association between stimulus B and the reward is not learned. This evidence suggests that learning is driven by the surprise; that is, by the difference between realized and expected value, the term PE.

The discovery of dopamine neurons (DN) has clarified the neural structure underlying learning by prediction error, and has provided an explanation of how the algorithm described in equations (2–5) is implemented. DN synthesize and release a neurotransmitter, dopamine. They are located in the ventro-anterior midbrain (substantia nigra and ventral tegmental area) and project to the striatum, ventral striatum, and nucleus accumbens, all structures involved in the evaluation of rewards. The DN respond with a short, phasic response (that is, responsive to change in stimulus, not to its sustained presentation).

The key finding (Tomas Ljungberg, Paul Apicella, and Wolfram Schultz 1991, 1992) is that the activity of the DN is proportional to the difference between the realized and the expected value of the reward. Suppose that the subject, a thirsty monkey, has learned in repeated trials that a visual cue announces the delivery of a reward in the form of a squirt of juice that can be large or small with equal probability. Since the animal has been in the experiment repeatedly, it knows the expected value $E(X)$ of the random variable X describing the delivery of juice.

In first approximation (we will see later an important and interesting modification of this relation), at the moment in which the juice is delivered, the monkey observes a realized value x. The activity of the DN is proportional to the difference between the realized and the expected value:

$$a = \alpha(x - E(X)), \tag{6}$$

which is the core element in the algorithm described by equations (2) and (3). In summary, the idea that learning is based on prediction error, as a basic building block, seems soundly established.

For the more general model described by equations (2–5), we now have a wealth of results supporting the view that adaptive learning is well described by it. Recent results suggest that the striatum contributes in its two anatomical components (ventral and dorsal striatum), with the former corresponding to the critic and the latter corresponding to the actor (see, for example, John O'Doherty et al. 2004). The model has important applications. For example, it throws light on the problem of addiction, as the research of Barry J. Everitt and Trevor W. Robbins (2005) has shown. The research in this area is very large and active. The progress to date is reviewed, for instance, in Nathaniel D. Daw and Kenji Doya (2006); Doya (2008); O'Doherty, Alan N. Hampton, and Hackjin Kim (2007); and Adam Johnson, Matthijs A. A. van der Meer, and A. David Redish 2007). Research in the last few years has extended these results from rodents and primates, where they were first documented, to humans. See Peter Dayan

This content downloaded on Fri, 4 Jan 2013 15:03:54 PM
All use subject to JSTOR Terms and Conditions

and Yael Niv (2008) for a review of these recent developments, as well as a list of intriguing open problems.

III. Adaptive Coding

Does this model help to predict behavior? Let us consider an application. Bernheim's favorite instance of an axiom is the weak axiom of revealed preferences. Just as the model describing behavior in an inter-temporal problem as the optimal solution of a dynamic programming problem, the assumption is considered unrealistic, and economists are probably willing to agree that context matters. The problem is to determine how it matters. We will consider an example to show how the new method can throw light on this question, and how the model and the results we have reported can shed light on an issue that already has been discussed in decision theory. Our choices, and the revealed preferences, can depend on the context, or on the environment. This may introduce inconsistencies in our behavior that seem to contradict standard assumptions. Let us start from a classical example of this dependence.

Errors in Decision Making.—In his foundational book, Leonard J. Savage (Savage 1972), was aware of possible inconsistencies between the theory he was outlining and real human behavior. An interesting example of these inconsistencies is the following "error in decision making." Keep in mind that, of course, the prices are 1954 prices:

> A man buying a car for $2,134.56 is tempted to order it with a radio installed, which will bring the total price to $2,228.41, feeling that the difference is trifling. But, when he reflects that, if he already had the car, he certainly would not spend $93.85 for a radio for it, he realizes that he has made an error. (Savage 1972, 103)

The example is compelling, and a plausible explanation seems natural. In this example, the decision maker is deciding in two different environments, and the range of monetary values that are being considered varies widely between the two. While he buys a car, the range of values he is considering is in the order of thousands of dollars, and the tradeoff between money and a radio is seen from this point of view. When he buys the radio alone, the corresponding range is in the order of a few dollars. In the second environment, the opportunity cost of money is more finely evaluated, so the comparison with the utility of the radio is more precise.

Response to Environment Changes.—The previous model of reward learning allows us to understand this error as the outcome of a design that is optimal under constraints. To see this, we need to go back to the monkey evaluating the random squirts of juice. Suppose that the animal has been evaluating the prediction error according to equation (6) for a while, with the random amounts varying in a fixed range. Suppose, now, that this range changes. A new environment or a new experimental condition occurs. For example, the squirts may now be the realization of a random variable Y, with values scaled up by a factor of ten with respect to X. The signal provided by a is an internal signal that is provided for the correct learning

This content downloaded on Fri, 4 Jan 2013 15:03:54 PM
All use subject to JSTOR Terms and Conditions

of the value of the state, in the process described in equation (3). For physiological reasons, the range of this internal signal is limited. It is bounded below because no negative firing rate is possible for the DN, and is bounded above by the physiological limits to the firing rate.

This constraint forces us to give a closer look to the linear form assumed in (6). This form can only be approximately true. For extreme values of the prediction error, the response to further increases is limited. Let us assume, for simplicity, that the response function is linear in the range of feasible responses, and flat when it is outside that range. We now consider the problem of choosing optimally the response function, if the objective is to maximize the informational value of this internal signal, given this constraint, and some information on the environment.

Optimal Adjustments.—Should the constant α in equation (6) be the same when the random amounts are realizations of X and of Y, if we want this internal signal to be accurate? If it is the same for X and Y, then in one of the two cases, the choice is not optimal. For example, if it is designed to match the range of the X variable, then, when Y is observed, many values fall out of the feasible response range. The optimal choice must depend on the environment.

Is there any evidence that it does? Experimental evidence in which the range of the values offered to the monkey is changed suggests that it does (see Philippe N. Tobler, Christopher D. Fiorillo, and Schultz 2005). More precisely, the correct equation when the animal is in the environment X is not equation (6), but instead

$$a = \alpha\left(\frac{x - E(X)}{SD(X)}\right), \tag{7}$$

where $SD(X)$ is the standard deviation of X. So the internal coding governing the evaluation of the prediction error adjusts to the environment the way an optimal designer would prescribe. The error that Savage (1972) pointed out in the example may be the consequence of an optimal adjustment on an internal learning mechanism. This error may be learned away, or compensated, but the essential point is that what seems a bias is an optimal adjustment to the environment.

Optimal Coding in Visual Perception.—Since relying on experimental evidence may not be satisfactory to all readers, let me point out that the adaptive coding is a widespread phenomenon. It was first discovered in the visual system, and a brief description of this finding may better illustrate its importance. After a visual stimulus has stimulated the retina, this input is translated by an internal coding into an output that is then evaluated. Just as in the case of economic choices, it is important to discriminate between rewards, and in this case it is important to discriminate between different stimulus intensities. In its natural environment, an animal will face a typical distribution of the stimuli, and it is natural to assume that it will adjust optimally to this distribution. Let μ be this distribution on the set of real numbers I of inputs, which we may assume to be the unit interval. A neuron carries this information downstream to an output R, also the unit interval, and may adjust in a

This content downloaded on Fri, 4 Jan 2013 15:03:54 PM
All use subject to JSTOR Terms and Conditions

nonlinear way to this stimulus. Call f the nonlinear adjustment. This f corresponds to the simple linear adjustment that we have assumed in (6).

An optimal coding would require that all response levels are used with the same frequency, so that the distribution of the output maximizes the entropy. To do this, the output should have, no matter what μ is, a uniform distribution on the output range.[1] Given μ, this requirement puts a very tight constraint of the function f, since it has to satisfy the equality

$$\mu(\{x : f(x) \in [a,b]\}) = U([a,b]) \tag{8}$$

for any interval $[a,b]$, where U is the uniform distribution over the interval. An easy computation shows that the function f that satisfies tequation (8) is the cumulative distribution function of the measure μ, that is,

$$f(x) = \mu([0,x]). \tag{9}$$

The function f in (9) is derived from an optimality condition on the transmission of information. The physiological response function can be measured and compared to the hypothesized f.

Here is a striking fundamental result in neural coding (conjectured in Horace B. Barlow 1961, and verified in Simon B. Laughlin 1981). The two functions match, so the real response function (operated by the neurons in the visual system of, for instance, a blowfly, as in Laughlin 1981) is the optimal adaptive coding to the environment. The conjecture that the same coding extends to the evaluation of rewards has found supporting evidence in single neuron recording study in the monkey (*Macaca Fascicularis*), see Tobler, Fiorillo, and Schultz (2005).

There are now three open problems. The first is how these findings extend to humans (as opposed to the monkey's) processing of rewards. The second (see Dayan and Niv 2008) is to identify the mechanism by which the downstream mechanisms can feedback the necessary information on the current range of the stimuli to the structures that are devoted to the adjustment of the sensitivity of the receptors. The most important problem, however, is to determine the speed and effectiveness of the adjustment to a new environmental condition. If the adjustment is too slow compared to the speed of the environmental change, then the decision maker will face a new environment with the old sensitivity. Errors in decision making may emerge as a result, as in the example presented by Savage (when a decision maker considers the purchase of a radio while buying a car, he may fail to adjust to the new, smaller range of values and underestimate the cost of a radio).

A recent literature in economics provides a foundation on evolutionary grounds of preferences that take into account relative values instead of simply absolute values, and that adjust to the range of the environment (see, for example, Arthur J. Robson 2001; Ken Binmore 1994, 25; Larry Samuelson 2004; Luis Rayo and Gary S. Becker

[1] A similar procedure is used in the optimal enhancement of digital pictures called *Histogram Equalization*, precisely because it operates a nonlinear transformation that maps the original histogram of shades of gray into the flat histogram of the uniform distribution.

This content downloaded on Fri, 4 Jan 2013 15:03:54 PM
All use subject to JSTOR Terms and Conditions

2007). This research provides a foundation for models of utility based on aspiration levels, as well as on the relative comparison with peers' outcome.

The neuroeconomics results we have seen put strong restrictions on the specific form of the dependence on how the happiness function (Rayo and Becker 2007) depends on the range of values in the typical environment. They also provide evidence that this is not just the form we might expect evaluations to have, but the form they really have. Finally, they also suggest that comparison of one's outcome with peers' outcome may have two separate motivations, and possibly two separate neural structures. One looks at the outcome of others simply to learn what the best action for an individual is. In this view, envy is simply the social equivalent of regret (see Nathalie Camille et al. 2004; Giorgio Coricelli et al. 2005). The other motivation looks at the comparison because it provides information on social ranking (see Nadège Bault, Coricelli, and Rustichini 2008). An additional insight from the neuroscience and neuroeconomic research is that the neural structures for comparisons (with the outcome of the non chosen actions, as in regret, as well as with the outcome of the actions chosen by others, as in envy) seem common with at least nonhuman primates. So the specific reliance on the hunter-gatherer society to explain these features of learning and evaluations in humans may have to solve the problem of why such structures are shared.

IV. Conclusions

The methodological framework suggested by Bernheim (2009), and slightly reformulated here, seems to me new, and different from the one used in mainstream economics. It is also very ambitious, and it is for the moment more a target and a benchmark than a criterion that is readily applicable. In this respect, it is not different from the axiomatic method. It is easy to state the criterion that the behavior of an economic decision maker should be summarized by a small set of independent axioms that characterize choice as the output of a simple functional representation. It is much harder to provide such a system. So it should not be a surprise that the current research in neuroeconomics does not follow it in full. In addition to the fact that researchers in the field might not agree with the formulation given here, the objective difficulty of realizing all the steps in the program is a formidable obstacle.

Will neuroeconomics eventually achieve what it promises? Bernheim poses a challenge that can provide a measure of success or failure:

> Provide an example of a novel economic model derived originally from neuroeconomic research that improves our measurement of the causal relationship between a standard exogenous environmental condition—one with which economists have been historically concerned—and a standard economic choice (Bernheim 2009).

The challenge has several components. Some of them are harder. For example, it is not going to be easy to find a *novel* economic model, and this is a merit of the creativity of economic theory. Models addressing a large number of problems have not been missing. They are all very useful. For example, the algorithmic model presented in (3) is close to several of the adaptive learning models that have been

This content downloaded on Fri, 4 Jan 2013 15:03:54 PM
All use subject to JSTOR Terms and Conditions

proposed in the past (Dean P. Foster and Rakesh V. Vohra 1999; Sergiu Hart 2005; Hart and Andreu Mas-Colell 2000).

This creativity is also the current weakness of economic theory, particularly in the area of decision theory. We do not lack models, we lack the ability to reject some of them on the basis of a well formulated strategy research. Neuroeconomics can be a valuable tool in this direction. For example, there is now a wide agreement in the neuroscience community that the adaptive learning in animals has the form based of the prediction error, as described by the equations (2) and (3). These findings put a very strong restriction on the type of adaptive learning that is really operative in animal learning. An additional fundamental role of neuroeconomics is, in my view, that of pruning the multiplicity of models, and to make them closer to the hard experimental test we suggested.

REFERENCES

Allais, M. 1953. "Le comportement de l'homme rationnel devant le risque: critique des postulats et axiomes de l'école Américaine." *Econometrica*, 21: 503–546.

Barlow, Horace B. 1961. "Possible Principles Underlying the Transformations of Sensory Messages." In *Sensory Communication*, ed. Walter A. Rosenblith. Cambridge, MA: MIT Press.

Bault, Nadège, Giorgio Coricelli, and Aldo Rustichini. 2008. "Interdependent Utilities: How Social Ranking Affects Choice Behavior." *PLoS ONE*, 3(10): 1–10.

Bernheim, B. Douglas. 2009. "On the Potential of Neuroeconomics: A Critical (but Hopeful) Appraisal." *American Economic Journal: Microeconomics*, 1(2): 1–41.

Binmore, Ken. 1994. *Game Theory and the Social Contract: Playing Fair.* Vol. 1. Cambridge, MA: MIT Press.

Camille, Nathalie, Giorgio Coricelli, Jerome Sallet, Pascale Pradat-Diehl, Jean-René Duhamel, and Angela Sirigu. 2004. "The Involvement of the Orbitofrontal Cortex in the Experience of Regret." *Science*, 304(5674): 1167–70.

Coricelli, Giorgio, Hugo D. Critchley, Mateus Joffily, John P. O'Doherty, Angela Sirigu, and Raymond J. Dolan. 2005. "Regret and its Avoidance: A Neuroimaging Study of Choice Behavior." *Nature Neuroscience*, 8(9): 1255–62.

Dayan, Peter, and Yael Niv. 2008. "Reinforcement Learning: The Good, The Bad and the Ugly." *Current Opinion in Neurobiology*, 18(2): 185–96.

Daw, Nathaniel D., and Kenji Doya. 2006. "The Computational Neurobiology of Learning and Reward." *Current Opinion in Neurobiology*, 16(2): 199–204.

Doya, Kenji. 2008. "Modulators of Decision Making." *Nature Neuroscience*, 11(4): 410–16.

Ellsberg, Daniel. 1961. "Risk, Ambiguity, and the Savage Axioms," *Quarterly Journal of Economics*, 75(4): 643–69.

Everitt, Barry J., and Trevor W. Robbins. 2005. "Neural Systems of Reinforcement for Drug Addiction: From Actions to Habits to Compulsion." *Nature Neuroscience*, 8(11): 1481–89.

Foster, Dean P., and Rakesh Vohra. 1999. "Regret in the On-Line Decision Problem." *Games and Economic Behavior*, 29(1–2): 7–35.

Glimcher, Paul W., and Aldo Rustichini. 2004. "Neuroeconomics: The Consilience of Brain and Decision." *Science*, 306(5695): 447–52.

Gul, Faruk, and Wolfgang Pesendorfer. 2005. "The Case for Mindless Economics. " Unpublished.

Hart, Sergiu. 2005. "Adaptive Heuristics." *Econometrica*, 73(5): 1401–30.

Hart, Sergiu, and Andreu Mas-Colell. 2000. "A Simple Adaptive Procedure Leading to Correlated Equilibrium." *Econometrica*, 68(5): 1127–50.

Johnson, Adam, Matthijs A. A. van der Meer, and A. David Redish. 2007. "Integrating Hippocampus and Striatum in Decision-Making." *Current Opinion in Neurobiology*, 17(6): 692–97.

Kamin, Leon J. 1969. "Selective Association and Conditioning." In *Fundamental Issues in Associative Learning*, ed. N. J. Mackintosh and W. K. Honig, 42–64. Halifax, Canada: Dalhousie University Press.

Laughlin, Simon B. 1981. "A Simple Coding Procedure Enhances a Neuron's Information Capacity." *Zeitschrift für Naturforschung C*, 36(9–10): 910–12.

This content downloaded on Fri, 4 Jan 2013 15:03:54 PM
All use subject to JSTOR Terms and Conditions

Ljungberg Tomas, Paul Apicella, and Wolfram Schultz. 1991. "Responses of Monkey Midbrain Dopamine Neurons during Delayed Alternation Performance." *Brain Research*, 567(2): 337–41.

Ljungberg Tomas, Paul Apicella, and Wolfram Schultz. 1992. "Responses of Monkey Midbrain Dopamine Neurons During Learning of Behavioral Reactions." *Journal of Neurophysiology*, 67(1): 145–63.

O'Doherty, John, Peter Dayan, Johannes Schultz, Ralf Deichmann, Karl Friston, and Raymond J. Dolan. 2004. "Dissociable Roles of Ventral and Dorsal Striatum in Instrumental Conditioning." *Science*, 304(5669): 452–54.

O'Doherty, John P., Alan N. Hampton, and Hackjin Kim. 2007. "Model-Based fMRI and its Application to Reward Learning and Decision-Making." *Annals of the New York Academy of Sciences*, 1104: 35–53.

Padoa-Schioppa, Camillo, and John A. Assad. 2006. "Neurons in the Orbitofrontal Cortex Encode Economic Value." *Nature*, 441(7090): 223–26.

Rayo, Luis, and Gary S. Becker. 2007. "Evolutionary Efficiency and Happiness." *Journal of Political Economy*, 115(2): 302–37.

Robson, Arthur J. 2001. "Why Would Nature Give Individuals Utility Functions?" *Journal of Political Economy*, 109(4): 900–14.

Samuelson, Larry. 2004. "Information-Based Relative Consumption Effects." *Econometrica*, 72(1): 93–118.

Savage, Leonard J. 1972. *The Foundation of Statistics*. New York: Dover Publications Inc.

Schultz, Wolfram, Peter Dayan, and P. Read Montague. 1997. "A Neural Substrate of Prediction and Reward." *Science*, 275(5306): 1593–99.

Schultz, Wolfram. 2002. "Getting Formal with Dopamine and Reward." *Neuron*, 36(2): 241–63.

Simon, Herbert A. 1957. "A Behavioral Model of Rational Choice." In *Models of Man, Social and Rational: Mathematical Essays on Rational Human Behavior in a Social Setting*. New York: John Wiley & Sons.

Sutton, Richard S., and Andrew G. Barto. 1998. *Reinforcement Learning: An Introduction*. Cambridge, MA: MIT Press.

Tobler, Philippe N., Christopher D. Fiorillo, and Wolfram Schultz. 2005. "Adaptive Coding of Reward Value by Dopamine Neurons." *Science*, 307(5715): 1642–45.

This content downloaded on Fri, 4 Jan 2013 15:03:54 PM
All use subject to JSTOR Terms and Conditions