A Model of Observational Learning with Boundedly Rational "Neuro" Agents

Kfir Eliaz

Tel Aviv University and University of Michigan

Ariel Rubinstein

Tel Aviv University and New York University

Abstract. We propose a simple model in which a boundedly rational agent observes not only the choices made by others, but also some information about the process that led them to those choices. We consider two cases: In the first, an agent observes whether another agent has compared the alternatives before making his choice. In the second, he also observes whether the decision was hasty. It is shown that the probability of making a mistake is higher in the second case and that the existence of these non-standard "neuro" observations systematically biases the equilibrium distribution of choices.

Keywords: Bounded Rationality, Neuroeconomics, Observational Learning.

The second author acknowledges financial support from ERC grant 269143. We would like to thank Zvi Artstein, Hadar Binsky, Tom Cunningham and especially Neil Thakral for their assistance.

1. Introduction

The standard method of modeling a decision maker in economics utilizes the concept of a choice function, which assigns a single alternative ("the choice") to every subset of available alternatives ("a choice problem") in some relevant domain. Thus, $C(A) = a$ means that the agent chooses the alternative *a* from the set *A*. Recent advances in choice theory have extended the traditional definition of a *choice problem* to include additional information, referred to as a *frame.* A frame represents the circumstances in which the choice problem was encountered, circumstances that do not affect the preferences of the decision maker, but may nevertheless affect his choice. Thus, $C(A,f) = a$ means that the decision maker will choose *a* from the choice problem *A* when *A* is presented in terms of the frame *f* (see Salant and Rubinstein (2008)). Leading examples of frames include a default option, the order in which alternatives are presented and the language in which the problem is phrased.

In this paper, we extend the choice function in a different direction as suggested in Rubinstein (2008) (see also Caplin and Dean (2011)). Instead of enriching the description of the *input* into the choice function (i.e., the choice problem), we enrich its *output*. For every choice problem, our augmented choice function specifies not only what the agent will choose, but also *evidence of the process that leads him to that choice*. Thus, $C(A) = (a, e)$ means that when an agent faces a choice set *A* he chooses *a* and produces evidence *e*. Examples of such evidence include response time, physical responses, such as blushing, and brain activity. An agent who is described by such an extended choice function is referred to as a "neuro agent".

The reader may wonder about the use of the term "neuro agent". Neuro information is usually thought of as information obtained by measuring various activities in the brain. We take a broader view of the term "neuro evidence" that includes any potentially observable information that a decision maker generates while making a choice. We use the term "neuro agent" to emphasize that from an "economic" point of view, it makes no difference whether the information is obtained by placing the agent in an fMRI machine or whether it is obtained through more conventional methods.

The novelty of this paper is the embedding of neuro agents within an economic model, where neuro information affects the their decisions. For example, suppose you need a dentist while on a business trip. You meet two individuals in a similar situation who have each already chosen a dentist. One deliberated for a long time and visited several dentists before making his choice. The other used the first dentist he found without making any comparisons. It seems plausible that you would be inclined to adopt the choice of the first individual.

The agent in our model is looking for a good or service that will satisfy some need and has two options to choose from. In order to compare them, he must meet two agents who have chosen different options. For example, suppose you are looking for alternative treatments of back pain. You know that there are various treatments available, but you are not familiar with them, nor do you know who can provide them. . In order to find the treatment that best suits you, you need to meet someone who has received that treatment. Only after different individuals introduce you to the available methods and give you contact information will you be able to make the appropriate choice.

An agent's behavior in our model is described without explicitly specifying an optimization that produces it. To optimally process the stream of information that agents receive and to arrive at the correct inference requires highly sophisticated skills that are possessed by most decision makers (even in our simple set up). Therefore, we study exogenously-given choice rules which can be viewed as heuristics that agents can revert to when faced with the complicated inference problem described here.

In the benchmark model, each agent sequentially samples up to *n* randomly-drawn observations of agents who have already solved the same decision problem. As soon as an agent observes two others who have made different choices, he stops the search, compares the two options and makes his own choice. Otherwise, he chooses the only observed chosen alternative. We study two variants of the model in which an agent's observation of another agent's behavior includes "neuro" evidence. In the first, the "neuro" evidence consists of whether an observed agent compared two options before choosing or merely imitated another agent's choice. The agent stops the search before observing two agents who made different choices if he reaches the end of the sample or if he observes an agent who made his choice after comparing the two options. In the second model, the neuro evidence consists of whether the agent decided hastily as soon as he observed one agent or waited longer. In this case, the agent stops the search without observing two different choices if he reaches the end of the sample or if he observes an agent who deliberated more than one period before making his choice.

In what follows, we define and characterize the equilibria of the models. We show that in the presence of neuro information the proportion of agents who choose the more "popular" option (i.e. the option more likely to be chosen following a comparison) is larger than in the absence of such information.

2. The model

There is a continuum of agents, each of whom chooses one of two alternatives: a or b. An agent's choice is determined by a procedure, that sequentially samples (at random) observations on other agents. An observation is a pair (x,e) where $x \in \{a,b\}$ is the choice of the observed agent and e is evidence of the choice procedure that produced the choice x. We refer to e as "neuro" evidence and assume that it takes values from some set E. A choice procedure is characterized by three parameters: the specification of E; a stopping rule that specifies after which sequences of observations the agent stops sampling and makes a decision; and the decision rule that assigns a lottery over $\{a,b\}$ to every sequence of observations that leads to stopping the search. In all the variants of the model, the agent stops searching as soon as he observes two agents who chose different alternatives. If the agent did not stop beforehand, he compares the two alternatives and with probability θ_x finds x to be preferred. Assume $\theta_a = \theta > 1/2$.

The symbol π_x^e represents the proportion of agents in the population who choose x and generate e. Denote by π the vector of $(\pi_x^e)_{x,e}$. We use the notation $\pi_x = \sum_e \pi_x^e$ for the frequency of agents who choose x and $\pi^e = \sum_x \pi_x^e$ for the frequency of agents who produce the evidence e. A stopping rule and a distribution of observations π induces a distribution $P(\pi)$ of the observations produced by an agent who samples from π and applies the stopping rule.

A *neuro equilibrium* is defined as a distribution π^* for which $P(\pi^*) = \pi^*$, i.e., in equilibrium, the distribution of observations of "newcomers" is identical to that of the existing population.

In order to define the notion of stability, let Δ be the set of probability distributions over $X \times E$. In each of the three models analyzed below, we specify a set $\Delta^* \subseteq \Delta$, which contains the possible distributions of observations. The set must satisfy the condition that the dynamic system, defined by $\dot{\pi} = P(\pi) - \pi$, remains within Δ^* for every initial condition within Δ^*. We say that an equilibrium $\pi^* \in \Delta^*$ is *stable* if the dynamic system is Lyapunov stable at π^*. In other words, for every $\varepsilon > 0$ there exists a δ small enough that, if the system starts within distance δ from π^*, it remains within distance ε from π^*.

3. The Benchmark Model

In the benchmark model, an agent observes only the choices of other agents (formally, E is a singleton). We assume that he follows procedure **(S-n)**, according to which he sequentially samples up to n agents and stops sampling as soon as either (i) he has sampled two agents who have made different choices or (ii) he has sampled n agents who all made the same choice. In case (i), he makes a comparison and chooses a with probability θ and b with probability $1-\theta$. In case (ii), he chooses the only option he has observed. This procedure induces the following function:

$$P_a(\pi) = (\pi_a)^n + \theta(1-(\pi_a)^n - (1-\pi_a)^n).$$

Note that the model always has two degenerate equilibria in which all agents choose one particular alternative. We are interested in *interior equilibria* which are characterized by a non-degenerate mixture of alternatives. With respect to stability, we will not impose any constraints on the possible distributions, i.e., $\Delta^* = \Delta$.

Proposition 0.
(i) *If $n > \frac{1}{1-\theta}$, then there exists a unique interior neuro equilibrium which is the only stable equilibrium. In this equilibrium, $\pi_a > \theta$.*
(ii) *The interior equilibrium converges to $(\theta, 1-\theta)$ as $n \to \infty$.*
(iii) *If $n \leq \frac{1}{1-\theta}$, then there exist only extreme neuro equilibria and the unique stable equilibrium is the one concentrated on a.*

Proof. Since $\pi_a + \pi_b = 1,$ the dynamic system is captured by the following function g, which describes the a-component of the dynamic system:

$$\dot{\pi}_a = g(\pi_a) = (\pi_a)^n + \theta(1-(\pi_a)^n - (1-\pi_a)^n) - \pi_a$$

A distribution $(\pi_a, 1-\pi_a)$ is an equilibrium if and only if $g(\pi_a) = 0$.

Note that $g(0) = g(1) = 0$, $g'(0) > 0$ and $n > \frac{1}{1-\theta}$ if and only if $g'(1) > 0$. It is straightforward to verify that for $n > 2$, there exists a unique interior value of π_a at which $g''(\pi_a) = 0$ and that for $n = 2$ there is no such value.

(i) It follows from the above that the function g must have an interior root. Furthermore, there exists a unique interior equilibrium $\pi_a^* \in (0,1)$ since if there were more than one, then $g'(\pi_a)$ would have at least three interior roots and $g''(\pi_a)$ would have at least two. Furthermore, since $g(\theta) > 0$ (given $\theta > 1/2$), we conclude that $\pi_a^* > \theta$.

The stability of the unique interior equilibrium follows from the fact that $g(\pi_a)$ is positive for $\pi_a < \pi_a^*$ and negative for $\pi_a > \pi_a^*$. Since the derivative of g is positive at the extreme points, g is positive near zero and negative near one and therefore, the degenerate equilibria are unstable.

(ii) Index the function g as g_n. The sequence of functions g_n converges to the function $\theta - \pi_a$, which equals zero only at $\pi_a = \theta$. Therefore, the sequence of interior equilibria must converge to $(\theta, 1-\theta)$ as $n \to \infty$.

(iii) Recall that $g(0) = g(1) = 0$, $g'(0) > 0$ and $g'(1) \leq 0$. Since $g''(\pi_a)$ has at most one interior root, $g'(\pi_a)$ has at most two. But if an interior equilibrium did exist, then $g'(\pi_a)$ would have at least three interior roots. ■

Thus, when n is not too small, the equilibrium proportion of agents who choose a is greater than θ. However, the excess of a-choosers goes to 0 as n increases. Furthermore, only a very small fraction of agents will not make a comparison. It follows that the distribution of choices in the benchmark model is almost unbiased unless n is "very small". This will no longer be true when agents observe “neuro information” about the individuals they sample.

Our benchmark model is related to the literature on word-of-mouth and social learning, in which agents observe samples of other agents’ actions and decide which is best for them, based on their observations. In one line of research, each agent receives a noisy signal regarding his payoffs from a given set of options, which is correlated with the signal received by other agents. Each agent chooses his action optimally after having observed the actions of some other agents. Following Banerjee (1992) and Bikhchandani, Hirshleifer and Welch (1992), some of these models assume that agents arrive sequentially and that each one observes the actions of all his predecessors. In others, such as Banerjee (1993), each agent observes the payoffs and actions of only a sample of other agents.

A second line of research examines exogenously-specified rules of behavior, which are not derived as the solution to some optimization problem (most notable are Ellison and Fudenberg (1993,1995)). In these models, an agent decides between two alternatives in each period. He has a preferred alternative, but does not know which it is because payoffs are noisy. The information available to the agent consists of other agents’ payoffs, which are correlated with his own. In some of these models, an agent observes a summary statistic of past payoffs and chosen actions, while in others he observes a summary statistic of only the current period’s payoffs. For experimental evidence on heuristical observational learning, see Hohnisch et al. (2012).

4. Were the options compared?

Assume now that the agent observes not only the choice made by another agent, but also additional "neuro evidence", in this case whether or not the other agent compared the two alternatives before making his choice. Let $E = \{+, -\}$. The observation $(x, +)$ means that "he chose x and made a comparison" and the observation $(x, -)$ means that "he chose x and did not make a comparison". Let π_x^+ and π_x^- denote the proportions of agents choosing x and producing the neuro evidence + and –, respectively.

Denote by **(C-n)** the procedure according to which an agent sequentially samples up to n other agents. As soon as he has sampled two agents who have made *different* choices, he stops, compares the two options and makes a choice. After a sequence of observations, $((x, -), (x, -), \ldots, (x, -), (x, +))$, of at most length n or after sampling the observation $(x, -)$ n times, the agent stops and chooses x.

The procedure is a rule of thumb which is not derived from the solution to an optimization problem. It seems reasonable to assume that this heuristic might be used by an agent who is ignorant of θ_a and does not understand the equilibrium. His only information is that there are two alternatives and that his preferred alternative is correlated with the preferred alternative of another agent. Thus, once he has observed another agent who has compared the two options he stops sampling and relies on that agent's choice. He does not make any inference from observing a long chain of individuals who have merely imitated one another.

As mentioned, this procedure is not derived from the solution to an optimization problem. Rather, we motivate the stopping rule as follows: Comparing the two options is the only way to ascertain one's own preferences. However, in order to make a comparison, the agent must wait for the two alternatives to appear. This may be costly for the agent since both sampling and comparing the two options may consume mental and physical resources. Therefore, given the correlation between the agent's preferences and those of other agents, it may be optimal for the agent to stop sampling once he has observed another agent who has compared the two options. On the other hand, it may be sub-optimal to stop searching after observing an agent who has made a choice *without* comparing the two options himself since, among other reasons, the individual's choice may be the outcome of a long chain of individuals who merely imitated one another since starting from the “initial state of the world”. This becomes even more likely if in the background there are also “noise” agents (not modeled here explicitly) who simply choose at random without sampling any other agents and without making a comparison. Thus, observing an agent who actually compared the two options seems intuitively to be more informative than observing an agent who simply mimicked another agent’s choice.

Given the assumption that agents who make a comparison choose x with probability θ_x, we restrict the set of distributions of observations, Δ^*, to those for which $\pi_a^+/\pi_b^+ = \theta/(1-\theta)$.

The above procedure induces the following P function: for $x = a, b$,

$$P_{(x,-)}(\pi) = \sum_{l=0}^{n-1} (\pi_x^-)^l \pi_x^+ + (\pi_x^-)^n$$
$$P_{(x,+)}(\pi) = \theta_x(1 - P_{(a,-)}(\pi) - P_{(b,-)}(\pi))$$

Note that the dynamic system $\dot{\pi} = P(\pi) - \pi$ remains in Δ^* since $\sum_{x=a,b}[P_{(x,-)}(\pi) + P_{(x,+)}(\pi)] \equiv 1$ and $P_{(a,+)}(\pi)/P_{(b,+)}(\pi) \equiv \theta/(1-\theta)$.

For the case $n = \infty$, we define $P_{(x,-)}(\pi) = \pi_x^+/(1-\pi_x^-)$ at any point where $\pi_x^- < 1$ and $P_{(x,-)}(\pi) = 1$ if $\pi_x^- = 1$.

In what follows, we focus on the two extreme cases, $n = 2$ and $n = \infty$, for which we establish the uniqueness and stability of interior equilibria.

Proposition (C-2). *Let $n = 2$. For $\theta \geq 2/3$, there is no interior neuro equilibrium. For $1/2 < \theta < 2/3$, there exists a unique interior equilibrium, which is stable and in which the proportion of a-choosers is $3\theta - 1 > \theta$.*

Proof. In equilibrium,

$$\pi_x^- = \pi_x^+ + \pi_x^-(\pi_x^- + \pi_x^+)$$
$$\pi_x^+ = \theta_x \pi^+$$

for $x = a, b$. It follows from the first equation that $\pi_a^-(\pi_b^- + \pi_b^+) = \pi_a^+$ and $\pi_b^-(\pi_a^- + \pi_a^+) = \pi_b^+$. The two equations imply that $\pi_b^+(\pi_a^- + 1) = \pi_a^+(\pi_b^- + 1)$ and hence,

$$(\pi_a^- + 1)/(\pi_b^- + 1) = \theta/(1-\theta)$$

The left-hand side must be less than 2 and therefore θ must be less than 2/3. In other words, for $\theta \geq 2/3$ the only equilibria are the extreme ones.

Let $f(z) = \frac{z(1-z)}{1+z}$. Thus, $\pi_a^+ = f(\pi_a^-)$ and $\pi_b^+ = f(\pi_b^-)$. The existence of an equilibrium is equivalent to the existence of a solution to the following equation:

$$1 = (\pi_a^- + \pi_a^+) + (\pi_b^- + \pi_b^+) = h(\pi_a^-) + h((\pi_a^- + 1)(1-\theta)/\theta - 1)$$

where $h(z) = z + f(z) = \frac{2z}{1+z}$. Since h is increasing, there is at most one solution for π_a^-. It is straightforward to solve the equation (for $\theta < 2/3$) and verify that the following tuple is an equilibrium:

$$(\pi_a^-, \pi_b^-, \pi_a^+, \pi_b^+) = \left(\frac{3\theta - 1}{3(1-\theta)}, \frac{2 - 3\theta}{3\theta}, \frac{(3\theta-1)(2-3\theta)}{3(1-\theta)}, \frac{(3\theta-1)(2-3\theta)}{3\theta} \right)$$

In this equilibrium, $\pi_b = (2 - 3\theta)$ and $\pi_a = 3\theta - 1 > \theta$.

For stability, note that a point in Δ^* is characterized by two parameters, π_a^- and π_b^-. The dynamic system can therefore be written as ($x = a, b$):

$$\dot{\pi}_x^- = \theta_x(1 - \pi^-) + \pi_x^-(\pi_x^- + \theta_x(1-\pi^-)) - \pi_x^-$$

Its Jacobian in equilibrium is:

$$\begin{pmatrix} -1 - \theta\pi_b^- + 2\pi_a^-(1-\theta) = \frac{9\theta-7}{3} & -\theta(1+\pi_a^-) = -\frac{2\theta}{3(1-\theta)} \\ -(1-\theta)(1+\pi_b^-) = -\frac{2(1-\theta)}{3\theta} & -1 - (1-\theta)\pi_a^- + 2\pi_b^-\theta = \frac{-9\theta+2}{3} \end{pmatrix}$$

It is straightforward to verify that the eigenvalues of this matrix are negative in the relevant range of θ. Therefore, the interior equilibrium is Lyapunov stable. ■

The next result presents a sufficient condition for the existence of an interior equilibrium for every $n > 2$. We believe that the equilibrium in (C-n) is unique, stable and has the property that more than θ of the participants choose A, though we have not been able to prove this analytically.

Proposition (C-n). *If* $\theta < \frac{2(n-1)}{2n-1}$, *then an interior neuro equilibrium exists.*

Proof. Define

$$f(y) = \frac{(y - y^n)(1-y)}{1 - y^n} = \frac{y - y^n}{\sum_{k=0}^{n-1} y^k}$$

Note that $f(0) = f(1) = 0$, $f'(0) = 1$ and $f'(1) = \frac{1-n}{n}$. In equilibrium ($x = a, b$):

$$f(\pi_x^-) = \theta_x \pi^+$$
$$\pi_a^- + \pi_b^- + \pi^+ = 1$$

An interior equilibrium exists if and only if there exists a solution to the equation:

$$g(y) = f(1 - y - f(y)/\theta) - (1 - \theta)f(y)/\theta = 0$$

That is, y^* is a solution to the above equation if and only if in equilibrium, $\pi_a^- = y^*$, $\pi_a^+ = f(y^*)$, $\pi_b^- = 1 - y^* - f(y^*)/\theta$ and $\pi_b^+ = f(y^*)(1 - \theta)/\theta$.

Note that $g(0) = g(1) = 0$ and $g'(y) = f'(1 - y - f(y)/\theta)(-1 - f'(y)/\theta) - f'(y)(1 - \theta)/\theta$. Hence,

$g'(0) = \frac{(2n-1)\theta-1}{n\theta} > 0$ for all $\theta > 1/2$, and

$g'(1) = \frac{(n-1)(2-\theta)-n\theta}{n\theta} > 0 \quad \text{iff} \quad \theta < \frac{2(n-1)}{2n-1}$

It follows that if $\theta < \frac{2(n-1)}{2n-1}$, then there exists y^* satisfying $g(y^*) = 0$ and hence an interior equilibrium exists. ■

The next result analyzes the equilibrium for the procedure (C-∞) in which the agent stops searching only if he observes the two options or if he samples another agent who has compared them.

Proposition (C-∞). *For $n = \infty$, there is a unique interior neuro equilibrium and it is stable. In this equilibrium: (i) the proportion of a-choosers is larger than θ and smaller than the proportion of a-choosers in the interior equilibrium for $n = 2$, and (ii) the probability that an agent makes a wrong decision is* $\frac{1}{2} - \frac{1}{2}\sqrt{4\theta - 4\theta^2 + 1}\,) + 2\theta(1 - \theta)$.

Proof. An interior equilibrium satisfies the following equations ($x = a, b$):

$$\pi_x^- = \frac{\pi_x^+}{1 - \pi_x^-}$$

$$\pi_x^+ = \theta_x \pi^+$$

Therefore, $\pi_a^-(1 - \pi_a^-) - (1 - \pi^+ - \pi_a^-)(\pi^+ + \pi_a^-) = (2\theta - 1)\pi^+$. Since $\pi^+ \neq 0$, we obtain $\pi^+ = 2\theta - 2\pi_a^-$. Substituting this into the first equation yields $(\pi_a^-)^2 - \pi_a^-(1 + 2\theta) + 2\theta^2 = 0$. The only solution of this equation, which is less than one, is:

$$\pi_a^- = \left(\frac{1}{2} + \theta\right) - \frac{1}{2}\sqrt{4\theta - 4\theta^2 + 1}$$

(Note that $4\theta - 4\theta^2 + 1 > 0$ for all θ and since $\sqrt{4\theta - 4\theta^2 + 1} < 1 + 2\theta$ we have $\pi_a^- > 0$. For $\theta > 1/2$, we have $4\theta - 4\theta^2 + 1 > 2\theta - 1$ and thus, $\pi_a^- < 1$.) The proportion of a-choosers, $\pi_a^- + \theta\pi^+ = (\theta - \frac{1}{2})\sqrt{4\theta - 4\theta^2 + 1} + \frac{1}{2}$, is greater than θ and one can verify that it is less than $3\theta - 1$.

An agent of type *a* (*b*) makes a mistake whenever he chooses *b* (*a*) without making a comparison himself. It follows that the probability of making a mistake is $(1-\theta)\pi_a^- + \theta\pi_b^-$. Substituting the equilibrium values for π_a^- and π_b^- yields the expression in (*ii*).

With respect to stability, consider the following dynamic system:

$$\dot{\pi}_a^- = \frac{\theta(1-\pi_a^- - \pi_b^-)}{1-\pi_a^-} - \pi_a^-$$

$$\dot{\pi}_b^- = \frac{(1-\theta)(1-\pi_a^- - \pi_b^-)}{1-\pi_b^-} - \pi_b^-$$

The Jacobian is:

$$\begin{pmatrix} \frac{-\theta\pi_b^-}{(1-\pi_a^-)^2} - 1 & \frac{-\theta}{1-\pi_a^-} \\ \frac{-(1-\theta)}{1-\pi_b^-} & \frac{-(1-\theta)\pi_a^-}{(1-\pi_b^-)^2} - 1 \end{pmatrix}$$

We have verified that the eigenvalues at the equilibrium point are negative and hence the equilibrium is Lyapunov stable. ■

Thus, if agents observe whether other agents have made a comparison, the proportion of *a*-choosers exceeds the “natural level” of θ. Comparing the two extremes, the excess of *a*-choosers when there is no bound on the number of samples is smaller than in the case of only two samples. This suggests that the excess may decrease as the number of allowable samples increases. However, unlike the benchmark case, the excess of *a*-choosers remains positive even in the extreme case in which an agent may continue sampling ad infinitum. The probability that an agent makes a wrong decision depends of course on θ: its maximal value is $1 - 1/\sqrt{2} \approx 0.29$ and it decreases to zero as θ increases from 0.5 to 1.

5. Was the decision hasty?

Assume now that an agent can observe not only the choices of other agents, but also whether they deliberated over their decisions or made their choices hastily after observing only one other agent. Denote by (T-n) the procedure according to which an agent sequentially samples up to *n* observations. As soon as he observes two agents who have made different choices, he stops the search, compares the two options and chooses one of them. As before, he also stops searching once he has observed an individual who has searched for at least two periods. In this case, the agent makes the same choice as the observed agent. If he samples *n* individuals who made the same choice after searching for only one period, the agent stops the search and makes the same choice as they did.

Formally, $E = \{1, 2\}$. The observation $(x, 1)$ means that the sampled agent chose x "hastily", i.e., after only a *single* observation. The observation $(x, 2)$ describes an agent who chose x and sampled at least *two* other agents prior to his choice.

As previously, we do not derive the search procedure from the solution of an optimization problem. Agents are persuaded to choose an option x if they themselves have compared the two options and found x to be preferable or if they have observed another agent who chose x after some deliberation.

The (T-n) procedure induces the following function P $(x = a, b)$:

$$P_{(x,1)}(\boldsymbol{\pi}) = \pi_x^2$$
$$P_{(x,2)}(\boldsymbol{\pi}) = \left[\sum_{k=1}^{n-1}(\pi_x^1)^k\right] \cdot \pi_x^2 + \theta_x\left[(1-\pi_x)\sum_{k=1}^{n-1}(\pi_x^1)^k + (\pi_x)\sum_{k=1}^{n-1}(\pi_{-x}^1)^k\right] + (\pi_x^1)^n$$

The model always has two extreme equilibria in which all individuals choose x (either a or b): half of the population does so immediately and the other half does so at a later point in time.

We again are mainly interested in the interior equilibria. The following proposition establishes necessary and sufficient conditions for the existence of an interior equilibrium and proves that whenever such an equilibrium does exist, it is unique (though we have not proven that it is stable). As before, we will deal separately with the analytically more convenient case of $n = \infty$, for which we will prove stability and show that the equilibrium proportion of a-choosers exceeds θ.

Proposition (T-n). *There exists an interior neuro equilibrium if and only if* $2 - (1/2)^{n-2} > \frac{\theta}{1-\theta}$. *When an interior equilibrium exists, it is unique.*

Proof. The equilibrium conditions are $(x = a, b)$:

$$\pi_x^1 = \pi_x^2$$
$$\pi_x^2 = \frac{\pi_x^1(1-(\pi_x^1)^{n-1})(\pi_x^2 + \theta_x\pi_{-x}^2 + \theta_x\pi_{-x}^1)}{1-\pi_x^1} + \frac{\theta_x\pi_{-x}^1(1-(\pi_{-x}^1)^{n-1})(\pi_x^2 + \pi_x^1)}{1-\pi_{-x}^1} + (\pi_x^1)^n$$

Define $A \equiv \pi_a^1$ and $B \equiv \pi_b^1 = 1/2 - A$. The above equations then reduce to:

$$A = \frac{A(1-A^{n-1})(A+2\theta B)}{1-A} + \frac{\theta B(1-B^{n-1})2A}{1-B} + A^n$$

Thus, an interior equilibrium exists if and only if the following equation has a solution in $(0,1)$:

$$\frac{1-\theta}{\theta} \cdot \frac{1-A^{n-1}}{1-A} = \frac{1-(\frac{1}{2}-A)^{n-1}}{1-(\frac{1}{2}-A)}$$

Letting $g(z) \equiv \frac{1-z^{n-1}}{1-z}$, we can rewrite the equation as follows:

$$\frac{1-\theta}{\theta} g(A) = g(\frac{1}{2}-A)$$

where $A \in [0, \frac{1}{2}]$. Note that $g(A)$ increases with A while $g(\frac{1}{2}-A)$ decreases with A. This has two implications: First, if an interior solution does exist, it is unique. Second, an interior solution exists if and only if $g(\frac{1}{2}) = 2-(1/2)^{n-2} > \frac{\theta}{1-\theta}$. ■

It follows from the proposition that for $n = 2$ there exist only extreme neuro equilibria. We have not been able to prove analytically that the proportion of *a*-choosers is higher than θ at the interior equilibrium of (T-n). The case of (T-∞) is much easier to fully address. In particular, we show that observing whether an agent made the choice hastily or not biases the equilibrium in favor of *a*.

Proposition (T-∞). *When $n = \infty$, there exists an interior neuro equilibrium if and only if $\theta < 2/3$. When this inequality holds, the equilibrium is unique and stable (for $\Delta^* = \Delta$). Furthermore, (i) the proportion of a-choosers is higher than that in the case of (C-∞), which in turn is higher than θ, and (ii) the probability that an agent makes a wrong decision is $\frac{1}{3}$.*

Proof. The equilibrium equations are ($x = a,b$):

$$\pi_x^1 = \pi_x^2$$
$$\pi_x^2 = \left[\sum_{k=1}^{\infty}(\pi_x^1)^k\right] \cdot \pi_x^2 + \theta_x\left[(\pi_{-x}^2 + \pi_{-x}^1)\sum_{k=1}^{\infty}(\pi_x^1)^k + (\pi_x^2 + \pi_x^1)\sum_{k=1}^{\infty}(\pi_{-x}^1)^k\right]$$

Denoting $A \equiv \pi_a^1$ and $B = \pi_b^1 = 1/2 - A$, we obtain:

$$A = \frac{A^2 + \theta 2AB}{1-A} + \frac{\theta 2AB}{1-B}$$

This equation has an interior solution $A = \frac{3\theta-1}{2}$ if and only if $\theta < \frac{2}{3}$. In the interior equilibrium, the probability of choosing *a* is $3\theta - 1 > \theta$. Furthermore, one can verify that the proportion of *a*-choosers is larger for (T-∞) than for (C-∞).

An agent of type *a* (*b*) makes a wrong decision after observing an agent who chose *b* (*a*) with some delay and none of the previous agents he observed had chosen *a* (*b*). It follows that the expected probability of making a mistake is:

$$(1-\theta)\frac{\pi_a^1}{1-\pi_a^1}+\theta\frac{\pi_b^1}{1-\pi_b^1}$$

Substituting the equilibrium values, $\pi_a^1 = \frac{3\theta-1}{2}$ and $\pi_b^1 = \frac{1}{2} - \pi_a^1$, we obtain that the probability of making a mistake is constant and equal to $\frac{1}{3}$ for all $\theta < \frac{2}{3}$.

To establish stability, we used Mathematica to derive the closed form expressions (as functions of θ) for the eigenvalues of the Jacobian matrix at the unique interior equilibrium. Using numerical methods, we then verified that all eigenvalues are negative when $\theta < 2/3$. ■

To conclude, when the hastiness of an agent is observed, the proportion of agents who choose *a* exceeds the "natural value" of θ, even in the limiting case where there is no bound on the number of samples. This excess is larger than in the case of the (C-∞) procedure. Furthermore, when $\theta \geq 2/3$, no interior equilibrium exists and the stable equilibrium is one in which all agents choose *a*. Note that the probability that an agent makes a mistake in equilibrium is higher for (T-∞) than for (C-∞). When $\theta < 2/3,$ the probability of making a mistake is $\frac{1}{3}$ for (T-∞) and at most 0.29 for (C-∞). When $\theta \geq 1/3,$ the equilibrium probability of making a mistake in (T-∞) is $1-\theta$ (since all agents choose *a*), which is larger than $\frac{1}{2} - \frac{1}{2}\sqrt{4\theta - 4\theta^2 + 1}\,) + 2\theta(1-\theta)$, the probability of making a mistake in (C-∞), for all $2/3 \leq \theta \leq 1$.

6. Final Comments

Neuroeconomics can be broadly viewed as the study of non-standard data, which includes information not only on individuals' choices but also on the choice *processes* they employ. To illustrate the relevance of such rich data in economics, our paper demonstrates how observations of other agents' choice *processes* of other agents (which we refer to as "neuro" evidence) influences an agent's decisions and to what extent they influence economic interactions. This paper explores this question using a simple model in which agents decide between option *a* and *b* applying choice procedures that take neuro evidence on other agents' choice processes as input and in equilibrium interpret this evidence in a consistent manner. We show that the use of neuro evidence leads to a stable equilibrium in which the proportion of *a*-choosers is higher than the actual proportion of agents who prefer *a* to *b*. Since this is only one particular example of a neuro model, future research should examine more interesting classes of models in which neuro information plays a crucial role.

References

Banerjee, Abhijit (1992). “A Simple Model of Herd Behavior,” *Quarterly Journal of Economics* **107**, 797-817.

Banerjee, Abhijit (1993). “The Economics of Rumors,” *Review of Economic Studies* **9**, 309-327.

Bikhchandani, Sushil, David Hirshleifer and Ivo Welch (1992). “A Theory of Fads, Fashions, Custom, and Cultural Change as Informational Cascades,” *Journal of Political Economy* **100**, 992-1026.

Caplin, Andrew and Mark Dean (2011). “Search, Choice and Revealed Preference,” *Theoretical Economics* **6**, 19-48.

Ellison, Glenn and Drew Fudenberg (1993). “Rules of Thumb for Social Learning,” *Journal of Political Economy* **101**, 612-643.

Ellison, Glenn and Drew Fudenberg (1995). “Word-of-mouth Communication and Social Learning,” *Quarterly Journal of Economics* **110**, 93-126.

Hohnisch, M., S. Pittnauer, R. Selten, A. Pfingsten and J. Eraßmy (2012). “Effects of Social Context on Decisions under Profound Uncertainty: Non-Robustness of Gender Differences,” Working Paper, University of Bonn.

Rubinstein, Ariel (2008). “Comments on NeuroEconomics,” *Economics and Philosophy* **24**, 485-494.

Salant Yuval and Ariel Rubinstein (2008). “(A,f), Choice with Frames,” *Review of Economic Studies* **75**, 1287-1296.