“Neuro”-Observational Learning

Kfir Eliaz
Brown University and Tel Aviv University

Ariel Rubinstein
Tel Aviv University and New York University

Abstract. We propose a simple model in which an 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, an agent observes the time invested in deliberation before the other agent makes his choice. It is shown that the existence of these non-standard “neuro” observations systematically biases the equilibrium distribution of choices.

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

Our goal in this paper is to illustrate how a broad view of "neuroeconomics" can be incorporated into economic modeling. We propose a model with "neuro-agents", i.e., agents who observe information on the process by which other agents reach their decisions.

The standard scheme of modeling a decision-maker in economics is through 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 a *frame.* A frame represents the circumstances in which the choice problem was encountered - circumstances, which do not affect the preferences of the decision-maker, but may nevertheless impact his choice. Thus, $C(A,f) = a$ means that the decision-maker will choose a when facing the choice problem A, presented in terms of the frame f (see Rubinstein and Salant (2007)). 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 (see Rubinstein (2007)). Instead of enriching the description of the *input* to the choice function (i.e., the choice problem), we enrich its *output*. For every choice problem, our augmented choice function specifies what the agent will choose, as well as *evidence about the process that leads the agent 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 activities. An agent who is described by such an extended choice function is referred to as a "neuro agent". Thus, we take a broad view of the term "neuro evidence" to include any potentially observable information that a decision-maker generates while making a choice. Understanding neuro agents is the objective of Neuroeconomics. The novelty of this paper is the embedding of neuro agents within an economic model.

Brain researchers are interested in "neuro" evidence of choice processes for its own sakes. For an economic agent, such information can be instrumental in decision-making since it may be helpful in interpreting the observed actions of other agents. For example, suppose you are looking for a dentist, while on a business trip. You meet two individuals in a similar situation who have each chosen a dentist. One deliberated for a long time and visited several dentists before making his choice. The other in contrast, picked 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 agents in our model are looking for a good or service that will satisfy some need. They know there are two options out there, but don't know what they are. Only after they meet another agent who has already chosen one of the options, do they learn of that option. To be able to compare the two options, an agent needs to meet one agent who chose one of the options and another agent who chose the other. For example, suppose you are looking for alternative treatments for back pain. You know there are various methods of treatment available, but don't know what these are nor 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.

In the model, each agent randomly samples observations of how other agents solved this decision problem. As soon as an agent observes two others who made different choices, he stops the search, compares the two options and makes his own choice. However, an agent might also stop the search after learning about only one option, if his observation of the other agent includes "neuro" evidence that "persuades" him to trust that agent's choice. In the first of our models, this additional information consists of whether or not the observed agent compared two options before making his choice. In the second model, the neuro evidence consists of the time it took for the observed agent to decide (that is, how many observations the sampled agent made before deciding). Note that the agents' behavior in our model is described without explicitly specifying any optimization that produces it. Rationalization is possible in our opinion, but not necessary for the point being made in this paper.

In what follows, we define and characterize the equilibria of the model. We show that in the presence of neuro information, the proportion of agents who choose the more "popular" option (the option more likely to be chosen following a comparison) is larger than in the absence of such information. This suggests that if, for example, 70% of the agents chose *a* and only 30% chose *b* in an environment where neuro information is available, then we can conclude that fewer than 70% *actually* prefer *a* to *b*. In short, the model demonstrates that a world in which neuro information is observed is quite different from one in which only decisions are observed.

2. The model

There is a continuum of agents in the model. Each agent chooses one of two available alternatives: *a* and *b*. Each agent is aware of the existence of the two options, but not of their substance. He is able to compare the two alternatives only after he observes two individuals, one who has chosen *a* and another who has chosen *b*. This could describe a situation in which an agent does not know which option he would prefer before he obtains information about the two options from two experienced individuals. In particular, it describes a situation in which an individual knows that two alternatives exist, but does not have access to an alternative $x \in \{a, b\}$ without being referred by an individual who has already chosen x.

Conditional on the comparison, a proportion θ_x of the agents would prefer x to the other alternative. Denote $\theta_a = \theta$ and $\theta_b = 1 - \theta$ and assume that $\theta > 1/2$. Thus, if an agent learns that another agent has already compared the two options and chosen x, then he would infer that it is more likely that he himself would also choose x after making the same comparison.

According to the conventional approach, an agent observes only the decision of another agent. Here we assume that he also observes additional details about the other agent's choice process. We refer to this extra information as "neuro" evidence. We denote an extended observation by (x, e), where $x \in \{a, b\}$ and e takes values from some set E. In what follows, we consider various specifications for E. 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 proportions over all extended observations. We use the notation $\pi^e = \sum_x \pi_x^e$ and $\pi_x = \sum_e \pi_x^e$.

An agent samples a sequence of agents who have already made a choice. He obtains extended observations of the sampled agents. As soon as he has observed two different choices he compares the two options and chooses the option he prefers. However, he might also stop the sampling and make a choice before seeing both options. The conditions under which this will occur are specified in the following sections.

The stopping rule induces a function $P(\pi)_{(x,e)}$ which specifies the probability, given the distribution of observations π, that an agent who applies the choice procedure will choose x and produce the extra information e. Let $P(\pi) = (P(\pi)_{(x,e)})_{(x,e) \in X \times E}$ be the vector of probabilities of the extended observations produced by the choice procedure given π.

We define a *neuro equilibrium* as a steady state of the system, which is characterized by a distribution π^* for which $P(\pi^*) = \pi^*$. In equilibrium, the distribution of extended observations on "newcomers" is identical to that on 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^*. That is, for every $\varepsilon > 0$ there exists a δ small enough such 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 choice 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 searching when either (i) he has sampled two agents who have made different choices, or (ii) he has sampled n agents who all chose the same alternative. In case (i), he makes a comparison and chooses his preferred option. In case (ii), he chooses the only option he has observed. This procedure induces the function:

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

Note that the model always has two degenerate equilibria in which all agents choose one particular alternative. We are interested in the *interior equilibria* which are characterized by a non-degenerate mixture of alternatives. With respect to the notion of 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. In this equilibrium* $\pi_a > \theta$. *This interior equilibrium is the only stable equilibrium.*

(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 function g, which describes the a-component of the dynamic system:

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

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 easy 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 equilibrium. There exists a unique interior equilibrium $\pi_a^* \in (0,1)$ since if there were more than one interior equilibrium, then $g'(\pi_a) = 0$ for at least three interior values of π_a, and $g''(\pi_a) = 0$ would have at least two interior values of π_a. 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, 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 $\pi_a = \theta$. Thus, 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 $g''(\pi_a) = 0$ for at most one interior value, $g'(\pi_a) = 0$ for at most two interior values. But if there were an interior equilibrium, then $g'(\pi_a)$ would equal zero for at least three interior values. ■

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. Furthermore, only a very small fraction of a-choosers will choose a before observing both alternatives and making a comparison. It follows that for values of n that are not "too small", the distribution of choices in the benchmark model is almost unbiased. This will no longer be true when agents observe "neuro information" about the individuals they sample.

4. When comparison is observed

Assume now that when one agent observes another, he observes not only the choice made, but also additional "neuro evidence", i.e., whether or not the other agent compared the two alternatives before making a 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 fraction of agents choosing x and producing the neuro evidence + and –, respectively. Also, let $\pi^+ \equiv \pi_a^+ + \pi_b^+$ and similarly for π^-.

According to the procedure we analyze in this section, denoted by **(C-n)**, 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.

This procedure is not derived as a 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 made a choice *without* comparing the two options himself. One reason for this is that an individual's choice may be the outcome of a long chain of individuals who merely imitated one another since the "initial state of the world". This becomes even more likely if in the background there are also "noise" agents (not modeled explicitly) who simply choose at random without sampling any 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 be all distributions 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^* because $\sum_{x=a,b}[P_{(x,-)}(\pi) + P_{(x,+)}(\pi)] \equiv 1$ and $P_{(a,+)}(\pi)/P_{(b,+)}(\pi) \equiv \theta/(1-\theta)$.

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

In the analysis below, 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 the proportion of a-choosers in this equilibrium 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 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 conjecture that the equilibrium in (C-n) is unique, stable and has the property that more than θ of the participants choose A. However, 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 $\theta < \frac{2(n-1)}{2n-1}$, 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 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, the proportion of a-choosers is larger than θ and smaller than the proportion of a-choosers in the interior equilibrium for $n = 2$.*

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 get $\pi^+ = 2\theta - 2\pi_a^-$. Substituting this into the first equation we get $(\pi_a^-)^2 - \pi_a^-(1 + 2\theta) + 2\theta^2 = 0$. 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 $3\theta - 1$.

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, once agents have observed 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.

5. Observing deliberation time

Assume next that an agent can observe not only the choices of other agents, but also the duration of their search, i.e., how many other agents they sampled before making their decision. This additional piece of information may be viewed as a proxy for the length of the observed agent's deliberation.

As part of the **(T-n)** procedure, an agent sequentially samples up to n observations. As soon as he observes two agents who have made two distinct choices, he stops the search, compares the two options and chooses one of them. He also stops searching once he observes someone 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 the n agents.

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.

We motivate the procedure as follows (as previously, we do not derive the search procedure from 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 observed another agent who chose x after deliberating for a "long enough" period. To understand why two periods of deliberation is "long enough", consider the following: Suppose all agents follow the (T-n) procedure but stop and imitate other agents only if they have deliberated for at least s periods, where $s > 2$. Consider an agent who observes another agent making a choice after s' periods, where $s > s' > 1$. He will conclude that the agent has made his choice after comparing the two options. Hence, the choice of the sampled agent is more informative than the choice of an agent who made a decision after s periods. It therefore seems reasonable (especially if sampling and making a comparison are costly) that the agent would stop the search after s' periods as well.

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: 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 next 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 this 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 also observing the length of deliberation 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 $\Delta^* = \Delta$) and the proportion of a-choosers is higher than that in the case of (C-∞), which in turn is higher than θ.*

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 $B = \pi_b^1 = A - 1/2$, 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-∞).

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 length of deliberation is observed, the proportion of agents who choose a exceeds the "natural value" of θ, even in the limit 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.

6. Related literature

The innovation of this model lies in its extension of the notion of observable information to include evidence of the choice *process*, such as response time or whether a comparison was made. The benchmark model (in which only choice is observed) is related to the literature on word-of-mouth and social learning, in which agents observe samples of other agents' actions and decide which action 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 other models, such as Banerjee (1993), each agent observes the payoffs and actions of only a sample of other agents.

In a different type of model, agents follow 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. Each agent 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 the agent observes a summary statistic of only the current period's payoffs.

The distinction between a choice function $C(A) = a$, choice with frames $C(A, f) = a$ and choice with neuro information $C(A) = (a, m)$ was suggested in Rubinstein (2007). Caplin and Dean (2011) consider an extended choice function that conveys not only the final option chosen by the decision-maker, but also how his choice changes during the period of deliberation prior to making the final selection. The authors characterize necessary and sufficient conditions for when such rich data sets are consistent with common search procedures.

7. Conclusion

This paper is motivated by a challenging question: to what extent can the ability of agents to observe signals regarding the choice process of other agents (which we refer to as "neuro"information) affect the outcome of economic interactions. We examine this question in the context of economic modeling, by constructing and analyzing a simple model in which agents receive neuro information (created non-strategically during the decision process) and in equilibrium interpret it in a consistent manner. Agents choose between two options, where $\theta > \frac{1}{2}$ of them would prefer a over b if they had the opportunity to compare the two options. The availability of neuro information affects the outcome of the stable equilibrium such that, the proportion of a-choosers exceeds θ.

We have looked at only one example of a neuro-model. Future research should be aimed at introducing other classes of models in which neuro information plays a crucial role.

References

Banerjee, Abhijit (1993). "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.
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.