Econ Theory (2014) 57:515–528
DOI 10.1007/s00199-014-0834-y

RESEARCH ARTICLE

A model of boundedly rational “neuro” agents

Kfir Eliaz · Ariel Rubinstein

Received: 21 October 2013 / Accepted: 18 July 2014 / Published online: 21 August 2014

Abstract We consider a model in which each agent in a population chooses one of two options. Each agent does not know what the available options are and can choose an option only after observing another agent who has already chosen that option. In addition, the agents' preferences over the two options are correlated. An agent can either imitate an observed agent or wait until he meets two agents who made different choices, in which case he can compare their choices and choose accordingly. A novel feature of the model is that agents observe not only the choices made by others, but also some information about the process that led them to those choices. We study two cases: In the first, an agent notes whether the observed agent imitated others or whether he actually compared the available alternatives. In the second, an agent notes whether the observed agent's decision was hasty or not. It is shown that in equilibrium the probability of making a mistake is higher in the second case and that the existence of these nonstandard “neuro” observations systematically biases the equilibrium distribution of choices.

Keywords Bounded rationality · Neuroeconomics · Choice procedures · Choice process data

K. Eliaz (✉)· A. Rubinstein
Eitan Berglas School of Economics, Tel Aviv University,
P.O.B. 39040, 69978 Tel-Aviv, Israel
e-mail: kfire@post.tau.ac.il

K. Eliaz
Economics Department, University of Michigan,
611 Tappan Ave., Ann Arbor, MI 48109, USA

A. Rubinstein
Economics Department, New York University,
19 W. 4th Street, New York, NY 10012, USA

JEL Classification D03 · D87

1 Introduction

The standard method of modeling an agent in economics utilizes the concept of a choice function, which assigns a single alternative ("a choice") to every subset of available alternatives ("a choice problem") in some relevant domain. The expression $C(A) = a$ is used to state that the agent chooses the alternative a from the choice problem 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 agent, but may nevertheless affect his choice. The expression $C(A, f) = a$ states that the agent chooses 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, we enrich its *output*. For every choice problem, our augmented choice function specifies not only what the agent chooses, but also *evidence of the process that leads him to that choice*. The expression $C(A) = (a, e)$ means that when an agent faces a choice problem A, he chooses a and produces evidence e. Examples of such evidence include response time, 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". Neuroinformation 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 the paper is that it embeds neuroagents within an economic model, in which neuroinformation affects their decisions. A typical situation we have in mind is the following: You need a dentist while on a business trip. You meet two individuals in a similar situation who have each already chosen different dentists. In addition to their choices you also observe information about their choice process. One individual deliberated for a long time while the other made a quick decision. It is likely 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. He knows that there are two options to choose from but he does not know what an option is until he meets some other agent who has chosen that option. Thus, he is able to compare the two options only after meeting two agents who made different choices. In addition, the agents' preferences over the two options are correlated. To illustrate this idea, suppose that you are looking for a treatment for 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 find someone who has had that treatment. As long as you meet only people who have made the same choice you can only imitate them. Only after you have met individuals who can inform you about both treatments and who can 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 the optimization that produces it. To optimally process the stream of information that agents receive and to arrive at the correct inference require highly sophisticated skills (even in our simple setup). 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 the observed agent has compared two options before making his choice or whether he 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 neuroevidence consists of whether or not the agent decided hastily as soon as he observed only one agent. 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 neuroinformation, the proportion of agents who choose the more "popular" option (i.e., the option more likely to be chosen following a comparison) is higher than in the absence of such information. Furthermore, this proportion is higher in the case when the neuroevidence includes whether or not a decision was made hastily. All proofs are presented in the appendix.

2 The model

There is a continuum of agents who wish to choose a good or service. Each agent has very limited knowledge of the available options. He knows that there are only two feasible alternatives, denoted as a and b, but does not know what they are. He cannot choose or even evaluate an alternative unless he meets an individual who has already chosen one of the alternatives. He knows that one of the two alternatives is the popular one in the sense that if he could compare the two options, he would prefer it with probability $\theta > 1/2$. However, he does not know which of the two alternatives it is, and he does not have any prior belief over the identity of the popular option. He understands that his preferred alternative is correlated with that of every other agent. Therefore, he perceives the choice made by another agent as being informative about the identity of the popular option. Without loss of generality, we denote a to be the popular alternative.

Ideally, the agent should compare the two alternatives himself and choose the preferred one. However, this requires him to wait until he meets two other agents who have made different choices. Alternatively, he can avoid the wait and exploit the correlation in preferences by imitating another agent's choice in the hope that he chose the preferred alternative. However, if that other agent did not compare the two alternatives himself but rather imitated another agent, then imitation may be less informative. This would be true if, for example, there was initially an arbitrary distribution of agents' choices and most agents just imitate one another without making a comparison.

An agent's choice in the model is determined by a *procedure* that sequentially samples (at random) observations of other agents. An observation is a pair (x, e) where $x \in \{a, b\}$ is the choice of the observed agent and $e \in E$ is evidence of the choice procedure that produced the choice x. We refer to e as "neuro" evidence.

Because the agents in the model have very limited knowledge, it seems more natural to assert that they follow a reasonable heuristic rather than assume that they solve a full-blown optimization problem (even though this latter approach is the more conventional one). A choice procedure in our model is a stopping rule that specifies the sequences of observations following which the agent stops sampling and makes a decision. In all the variants of the model, the agent stops searching as soon as he observes two agents who chose different alternatives. The variants differ in their specification of the stopping rule in that case that he only meets agents who made the same choice (with possibly different "neuro" information).

Our objective is to investigate the choice dynamics of a population of agents who follow a given choice procedure. The symbol π_x^e represents the proportion of agents in the population who choose x and generate the evidence e. Denote by π the vector $(\pi_x^e)_{x,e}$. 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 induce a distribution $P(\pi)$ of the observations produced by an agent who samples from π and applies the stopping rule.

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

In order to define the notion of stability, we need to specify a set Δ^* of possible distributions of observations. This 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^*. In two variants of the model, we set $\Delta^* = \Delta$, the set of probability distributions over $X \times E$. 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 leads to the following function:

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

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 nondegenerate 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 neuroequilibrium, 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 neuroequilibria, and the unique stable equilibrium is the one concentrated on a.

When n is large enough, most agents will eventually compare the two options, so that the distribution of choices will converge to the distribution of preferences in which a proportion of θ chooses a. If n is small, then in equilibrium, agents tend to imitate one another and hence will be less likely to compare the two alternatives themselves. In each of the two extreme equilibria, all agents choose the same option, but in the unique stable equilibrium, all agents choose a such that agents who would have preferred option b end up with the "wrong" option. Thus, unless n is sufficiently small, the equilibrium distribution of choices in the benchmark model is unbiased. This will no longer be true when agents observe "neuroinformation" about the individuals they sample.

The benchmark model is related to the word-of-mouth and the social learning literature in which agents observe samples of other agents' actions and then decide which is best for them. 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 et al. (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. In contrast to Proposition 0, these papers show that as the population of agents grows, the equilibrium converges to an inefficient outcome.

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. These models yield a result similar to the equilibrium characterization of our benchmark model: Despite the fact that agents are not optimizing, players may eventually adopt the action, which is on average superior. For experimental evidence on heuristical learning, see Hohnisch et al. (2013).

4 Were the options compared?

Assume now that the agent observes not only the choice made by another agent, but also additional "neuroevidence", 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" while 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 neuroevidence $+$ 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.

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 (especially if θ is large), it may be optimal for the agent to stop sampling once he has observed an agent who has compared the two options. However, it may not be optimal to stop searching after observing an agent who has made a choice *without* having compared the two options himself. This is because that agent's choice may be the outcome of meeting a long chain of agents who merely imitated one another starting from an arbitrary initial distribution. The above procedure also seems reasonable if in the background there are "noise" agents (not modeled here explicitly) who simply choose randomly without sampling any other agents and without making a comparison.

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 leads to 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 neuroequilibrium. For $1/2 < \theta < 2/3$, there exists a unique interior equilibrium, which is stable and in which the proportion of a-choosers is $3\theta - 1 > \theta$.*

The next result presents a sufficient condition for the existence of an interior equilibrium for every $n > 2$. While we have not been able to prove this analytically, we believe that this equilibrium has the following properties: (i) It is the only interior equilibrium in (C-n), (ii) it is stable and (iii) more than θ of the participants choose A.

Proposition (C-n.) *If $\theta < \frac{2(n-1)}{2n-1}$, then an interior neuroequilibrium 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 and stable interior neuroequilibrium. 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 $\frac{1}{2} - \frac{1}{2}\sqrt{4\theta - 4\theta^2 + 1}) + 2\theta(1-\theta)$, which is larger than $1 - \theta$.*

To summarize, this section has analyzed the case in which agents observe not only the choice of other agents they meet, but also whether that choice was the result of a comparison. Agents stop their search as soon as they meet an agent who has compared the two options himself.

When the number of search periods is small (two) and correlation is high, the system will settle on the extreme distribution where the "most popular" alternative (i.e. a) will be chosen by all agents, yielding a probability of mistake of $1 - \theta$. However, when the number of search periods increases, there exists an interior equilibrium. Unlike the benchmark case, in this equilibrium, the excess of a-choosers remains positive even in the extreme case in which an agent may continue sampling ad infinitum. Furthermore, the proportion of a -choosers in the limit exceeds the "natural level" of θ, and the probability of making a mistake exceeds $1 - \theta$. The precise probability of making a mistake depends on the value of θ but its maximal value is $1 - 1/\sqrt{2} \approx 0.29$. The probability of mistake is higher than if all agents were to simply choose a. In contrast, in the benchmark model, there are no mistakes in the limit. However, a welfare comparison between the two procedures is problematic since it ignores the benefits from shortening the search time.

5 Was the decision hasty?

Assume now that an agent can observe not only other agents' choices, 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. 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.

Note the relation between the neuroevidence observed in procedure (C-n) and that observed in (T-n). In (C-n), an agent might know for certain that another agent has compared the two options, while in (T-n), he might know for certain that another agent (who decided after only one period) has *not* compared the two options.

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 before, 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 leads to the following function P $(x = a, b)$:

$$P_{(x,1)}(\pi) = \pi_x^2$$

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

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

It follows from the proposition that for $n = 2$, there exist only extreme neuroequilibria. 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 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 neuroequilibrium 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* $\frac{1}{3}$.

In conclusion, when the hastiness of an agent can be 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, and this excess is larger than in the case of (C-∞). 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 $\theta \geq 2/3$, 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 broadly be viewed as the study of nonstandard data, which includes information not only on individuals' choices but also on their choice *processes* they employ. To illustrate the relevance of such rich data to economics, our paper demonstrates how the observations of other agents' choice *processes* (which we refer to as "neuro" evidence) influence an agent's decisions and the extent to which they influence economic interactions. A simple model was used in which agents decide between two options (a and b) by applying choice procedures that take neuroevidence as input and in equilibrium interpret this evidence in a consistent manner. The use of neuroevidence 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 neuromodel, future research should examine more interesting classes of models in which neuroinformation plays a crucial role.

Acknowledgments 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.

Appendix

Proof of Proposition 0. 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:

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

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. □

Proof of Proposition (C-2.) 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 <2, and therefore, θ must be <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. □

Proof of Proposition (C-n.) 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

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

Proof of Proposition (C-∞). 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. □

Proof of Proposition (T-n.) 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:

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

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 $g(\frac{1}{2}) = 2 - (1/2)^{n-2} > \frac{\theta}{1-\theta}$. □

Proof of Proposition (T-∞). 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 $\theta < 2/3$. □

References

Banerjee, A.: A simple model of herd behavior. Q. J. Econ. **107**, 797–817 (1992)
Banerjee, A.: The economics of rumors. Rev. Econ. Stud. **9**, 309–327 (1993)
Bikhchandani, S., Hirshleifer, D., Welch, I.: A theory of fads, fashion, custom, and cultural change as informational cascades. J. Polit. Econ. **100**, 992–1026 (1992)
Caplin, A., Dean, M.: Search, choice and revealed preference. Theor. Econ. **6**, 19–48 (2011)
Ellison, G., Fudenberg, D.: Rules of thumb for social learning. J. Polit. Econ. **101**, 612–643 (1993)
Ellison, G., Fudenberg, D.: Word-of-mouth communication and social learning. Q. J. Econ. **110**, 93–126 (1995)

Hohnisch, M., Pittnauer, S., Selten, R., Pfingsten, A., Eraßmy, J.: Effects of Social Context on Decisions Under Profound Uncertainty: Non-Robustness of Gender Differences. University of Bonn, Bonn, Mimeo (2013)
Rubinstein, A.: Comments on neuroeconomics. Econ. Philos. **24**, 485–494 (2008)
Salant, Y., Rubinstein, A.: (A,f), Choice with frames. Rev. Econ. Stud. **75**, 1287–1296 (2008)