Learning under Coarse Thinking

Daniel M. de A. Barreto *

Abstract

In this paper I build a framework for modelling agents that can hold beliefs belonging to a coarse space rather than a continuum in the context of a two-armed bandit problem. I show that coarseness of beliefs does not imply a difference in optimal behavior compared to agents that think in a continuum, such that any difference in exploration among these agents has to be related with the way they interpret evidence. I then discuss broad classes of interpretation rules and analyze what they imply in the exploration-exploitation dilemma. More specifically, I show that coarse thinkers whose interpretation roughly follows that of a fine thinker can be thought of as agents that think in a continuum that is partitioned into categories such as “likely”, “unlikely” and so on. For such agents, the length of exploration will ultimately depend on how the cutoff belief of the two-armed bandit is perceived.

Keywords: Information, Learning, Two-Armed Bandit, Coarse Thinking

1. Introduction

Learning is a process of observation and interpretation of evidence. People are constantly facing uncertainties about aspects of reality that are relevant for their lives, and in order to navigate through these uncertainties they act based on their beliefs. Naturally, beliefs are not static, and they can always be subject to change upon observation of new evidence. New evidence is then interpreted, process by which we try to assess their true informational content, and a (possibly new) belief is adopted.

A sailor exploring an unknown sea will observe signals such as drifting land vegetation or the behavior of birds around him in order to learn about whether or not he is approaching

*Email address: daniel.martinsdealmeidabarreto@sciencespo.fr

I would like to thank Professor Eduardo Perez-Richet for his guidance and advice. I would also like to thank my friends for the stimulating discussions and my family for the unconditional support and love.

land. A job-interviewer, faced by a candidate she has never seen before, will interpret the candidate's responses and CV in order to form a belief about whether he is suited or not for the job. A hungry tourist wandering through the streets of Paris might look at the movement in different restaurants while trying to assess which one is better.

In all of those cases individuals observe signals in order to form beliefs about an uncertain state. The type of evidence they observe, however, is different in each scenario: the sailor interprets signals from nature while learning if there is land nearby, the job-interviewer evaluates information conceded by an agent that is himself interested in getting hired while assessing his suitability for the job, and the hungry tourist observes the actions taken by other (formerly) hungry tourists while evaluating the quality of a restaurant.

Such differences help shedding light into different aspects of the transmission and accumulation of information. The problem of the hungry tourists, for example, is one of aggregation of dispersed information: each tourist holds some private imperfect information about the quality of the restaurants, such that a tourist interested in choosing the best restaurant to have dinner will not rely solely on his piece of information, but also try to infer the private information of other tourists by observing which restaurant they choose. A given tourist choosing between dining at Le Piège or Chez Augustin, for example, might have imperfect private information favoring the latter. Still, observing enough people choosing Le Piège might lead him to optimally ignore his private information and just herd along. If he does so, his presence in this restaurant will serve as further evidence for other tourists even though it does not reflect his private information. Such informational cascades are studied in the Social Learning literature (Banerjee [1], Bikhchandani *et al.*[2], Acemoglu *et al.*[3]).

The problem of the job applicant and the interviewer, on the other hand, is one of strategic information transmission: the job applicant is informed about whether he is suited or not for the job, so that he will choose how to disclose information about himself so as to induce the interviewer to hire him. Such separation between possession of information and agency motivates many interesting works in the Cheap Talk (Crawford and Sobel [4], Farrell and Rabin [5]) and Bayesian Persuasion literature (Kamenica and Gentzkow [6], Bergemann and Morris [7]).

Finally, the problem of the sailor that chooses whether he should continue exploring an unknown sea or sail back home is one of strategic exploration: the sailor has to choose how to optimally allocate his resources (in this case, his time) into exploring and learning about an uncertain state (whether or not this sea has islands he can discover) or exploiting a certain alternative (going back to his homeland). This exploration-exploitation dilemma has been studied in a variety of different settings (Bolton and Harris [8], Keller *et al.* [9], Fryer and Harms [10]).

Such problems are traditionally studied in a setting where agents have prior beliefs about the unknown state and know the conditional distributions of signals, so that they are able to fully extract the information contained in the evidence they observe. Modifications of such models in which the agents' rationality is bounded, however, are increasingly common (Mullainathan *et al.* [11], Guarino and Jehiel [12], Bohren and Hauser [13]).

In this paper we propose such a modification in the context of the exploration-exploitation dilemma. Instead of assuming an agent with an infinite belief space, we think about an agent that thinks about the uncertainties he faces in terms of coarse categories. This has a clear impact on the way agents interpret information: while agents that think in a continuum can map their information to an infinite number of points, coarse thinkers have a finite number of beliefs with which they can represent such information.

For the adventurous sailor, observing drifting land vegetation, shorebirds flying or even a distinctive pattern of swells in the sea will serve as signals that there are islands around him. Conversely, not observing those is suggestive that such islands might not exist. Thinking about this uncertainty in terms of coarse categories such as "very likely", "likely", "unlikely" and so on instead of in a continuum will necessarily imply a difference in the way he interprets the signals he is faced, which could potentially affect the length of time he will choose to engage in exploration.

Categorical thinking has for long been subject to study in psychology (Macrae and Bodenhausen [14], Markman and Gentner [15]). In Murphy and Ross [16], authors present an experiment in which people were exposed to different sets of drawings (containing geometrical shapes, faces, etc), each identified as having been drawn by a particular child. They were then shown a new drawing and asked which child they thought had drawn it and whether or not they thought the drawing would exhibit a given property (e.g. have originally been drawn in red). Whereas bayesian inference would require subjects to evaluate the recurrence of this given property over the drawings of every children, the study showed that subjects generally relied solely on the recurrence of this property in the category they used (e.g. drawings by Lucy). Further evidence that individuals often ignore information from alternative categories was found in Malt *et al.* [17], Kahneman and Tversky [18] and Kahneman, *et al.* [19].

Another interesting experiment is presented in Krueger and Clement [20]. In it, subjects where asked to estimate the average temperature in their city during several days in the year. The results exhibited greater differences in the predicted temperatures for days in different months than for days in the same month, even when the distance between the days in different months was smaller than the distance between the days in the same month. For example, the average prediction for the higher temperature in the 2nd of February was 38°F, which is very close to the average prediction for the higher temperature during the 24th of February: 42°F. The average prediction for the 4th of March (which is roughly a week away from the 24th of February), however, was significantly higher: 53°F. This suggests that the predictions relied almost entirely on the category of the day (i.e. the month).

Categorization naturally implies a notion of coarseness. Mullainathan [21] provides a model of inference under coarse thinking in which agents have their belief space partitioned into categories, such that every posterior categorized under a given category is perceived equally. For example, an agent could have its belief space over an unknown state $[0, 1]$ partitioned equally into 5 categories such that beliefs in $[0, 0.2)$ are treated as "very unlikely", beliefs in $[0.2, 0.4)$ are treated as "unlikely", and so on. This would imply a coarse set of adoptable beliefs such as M ={"very unlikely", "unlikely", "medium", "likely", "very likely"}, where each of these beliefs is associated to the posteriors that are categorized under them.

Our way of modelling coarseness of beliefs is more general. Although each of the beliefs in M ={"very unlikely", "unlikely", "medium", "likely", "very likely"} need to represent a point in $[0, 1]$, we do not constrain them to represent a specific partition on $[0, 1]$. A categorical representation might be possible depending on the interpretation rule adopted by the agent (we show the conditions for that), but that does not have to be necessarily the case.

The remaining of the paper is organized as follows: in Section 2 we introduce the two-armed bandit and then proceed to describe the coarse thinkers and to characterize their optimal behavior in the bandit problem. Section 3 discusses broad classes of interpretation rules and how they affect the exploratory behavior of agents and Section 4 concludes.

2. The Model

2.1. The Two-Armed Bandit

There is an infinite number of periods $t = 0, 1, ...$ and an unknown state of nature $\theta \in \{0, 1\}$. In each period an agent chooses an action $a_t \in \{R, S\}$, where the safe action $a_t = S$ yields a safe payoff normalized to 0 and the risky action $a_t = R$ yields a payoff $X_{i,t}$ that can be either x_L or x_H, with $x_L < 0 < x_H$. In the remaining of the paper we will sometimes refer to obtaining a payoff x_H as a success, and obtaining a payoff x_L as a failure.

The usual illustration of this problem is that of a gambler that has to choose how much resource (in this case time) to allocate into a slot machine that yields a known reward and into one whose payoff process is unknown to her and potentially better than the safe one, hence the Two-Armed Bandit name. By pulling the risky arm a payoff is realized, which then serves as evidence for an agent that is looking to learn about the true state of nature.

Risky payoffs are drawn independently at each period, conditional on the state of nature θ. When $\theta = 1$, the high payoff is given with probability $Pr(x_H|\theta = 1) = \pi \in (0, 1)$, whereas when $\theta = 0$ such payoff is never drawn (i.e. $Pr(x_H|\theta = 0) = 0$). Since in this model high payoffs x_H can only be drawn on the good state of nature, a single draw of x_H is fully revealing about θ.

We call E_j the conditional expectation $\mathbb{E}[X_{it}|\theta = j]$ for $j \in \{0, 1\}$. Since $x_L < 0$, we know that $E_0 = \mathbb{E}[X_{it}|\theta = 0] = x_L < 0$, and we assume that $E_1 = \mathbb{E}[X_{it}|\theta = 1] = \pi x_H + (1 - \pi)x_L > 0$.

Given a prior belief p_0 that the state of nature is good, the agent chooses an action plan $(a_t)_{t=0}^{\infty}$ in order to maximize the expected payoff

$$\max_{(a_t)_{t=0}^{\infty}} \quad \mathbb{E}_p[\sum_{t=0}^{\infty} \delta^t \mathbb{1}_{\{a_t=R\}} X_{i,t}] \tag{1}$$

where $\delta \in [0,1]$ is her time discount factor. This problem is often referred to as the exploration-exploitation dilemma, since it involves the trade-off between exploiting a certain alternative and gathering information about the process underlying an uncertain one.

Interpretation of evidence is done through belief updating. Optimal extraction of information from evidence requires that the belief is updated according to Bayes rule, that is, that after pulling the risky arm and observing payoff realization $X_{i,t} = x_L$ an agent with prior belief p forms a posterior belief given by $p' = \frac{(1-\pi)p}{p(1-\pi)+1-p}$. More generally, the belief of an agent after observing n failures and no success will be $B(n, p_0) = \frac{(1-\pi)^n p_0}{(1-\pi)^n p_0+1-p_0}$. Since payoff x_H can only be realized under $\theta = 1$, after the first success in the risky arm the belief is updated to $p = 1$.

By the dynamic programming principle we can focus on action plans in which a_t is a time-invariant function of the belief held at that moment, such that p is the relevant state variable. The value of a given state variable p is given by:

$$V(p) = max\{p\pi(x_H + \delta\frac{E_1}{1-\delta}) + (p(1-\pi) + (1-p))(x_L + \delta V(p')), 0\} \tag{2}$$

where $\frac{E_1}{1-\delta}$ is the value of exploring forever while knowing that the state is $\theta = 1$. When $V(p) = 0$ the agent exploits (i.e. plays S) and when $V(p) > 0$ she explores (i.e. plays R).

As such, for some values of p it will be optimal for the agent to play $a_t = R$, while for other values it will be optimal for the agent to play $a_t = S$. She will be indifferent between playing R or S for a given p for which $p\pi(x_H+\delta\frac{E_1}{1-\delta})+(p(1-\pi)+(1-p))(x_L+\delta V(p')) = 0$. We know in this scenario that for the next belief she will prefer playing S, such that $V(p') = 0$. The cut-off belief which makes her indifferent is then given by:

$$p^* = \frac{-(1-\delta)E_0}{(1-\delta)(E_1-E_0)+\delta\pi E_1} \tag{3}$$

An agent will thus choose $a_t = R$ as long as she holds a belief $p > p^*$. As soon as her belief reaches the cut-off p^* she will shift to exploitation on the safe arm S.

As an exercise we can consider an agent that only cares about the present period, i.e. an agent for whom $\delta = 0$. The cut-off belief for this myopic agent is then $p^m = \frac{-E_0}{E_1 - E_0}$, which is clearly greater than the previous cut-off p^*. This reflects the fact that a forward-looking agent values the information contained in the payoff realizations of the risky arm and will thus be willing to engage in exploration for longer than the myopic agent.

2.2. Coarse Thinkers

Coarse thinking is here treated as a bound on the belief space of agents. Instead of holding any belief on $[0, 1]$, coarse thinkers have a finite belief space $\mathfrak{M} = \{\bar{\mu}, \mu_0, \mu_1, ..., \mu_k, \underline{\mu}\}$ where $\mu_i \in (0, 1) \forall i \in \{0, 1, ..., k\}$. The only assumptions we make on this belief set are that (i) it includes certainty beliefs $\bar{\mu} = 1$ and $\underline{\mu} = 0$ and that (ii) all the other beliefs in the set are either the prior or beliefs smaller than the prior. Assumption (ii) is made without loss of generality since beliefs higher than the prior but different from 1 would never be adopted in this problem (as evidence here can only lower the agent's belief or make him certain that $\theta = 1$). For convenience we label the prior belief as μ_0, and the subsequent beliefs in the set in decreasing order.

We denote by $h_t = (y_0, y_1, ..., y_{t-1})$ the observable history of risky payoffs at time t, where $y_i \in \{x_L, x_H\} \forall i \in \{0, 1, ...\}$. An interpretation rule will then be a mapping $I : \{x_L, x_H\}^t \rightarrow \mathfrak{M}$.

We will here restrict our attention to mappings that satisfy the following basic criteria: (i) any history with at least one success is mapped into $\bar{\mu}$; conditional on not having observed any success so far, (ii) if a given history h_t is mapped into μ_i, then every other history h'_t with more failed attempts than h_t has to be mapped into a μ_j with $j \geq i$ and (iii) if a history h^l_t with l failed attempts is mapped into μ_i and a history h^n_t with n failed attempts is mapped into μ_{i+2}, then there exists a history h^m_t with $l < m < n$ failed attempts that is mapped into μ_{i+1}. Conditions (i) and (ii) ensure that we are working with interpretation rules that (correctly) see a success as conclusive evidence that $\theta = 1$ and failures as evidence suggestive that $\theta = 0$, while condition (iii) ensures that agents can potentially use all beliefs available for them.

Under such conditions the interpretation rules are going to define, for each belief μ_i, the number of failures to be observed before switching to belief μ_{i+1}. We will call $\bar{k}(\mu_i)$ the inertia of belief μ_i, which indicates the number of failures observed under belief μ_i before belief μ_{i+1} is adopted in a given interpretation rule. The interpretation rules will thus define a sequence $(\bar{k}(\mu_i))_{i=0}^{k}$ such that:

$$\mu(h_t) = \mu(f(h_t), s(h_t)) = \begin{cases} \bar{\mu}, & \text{if } s(h_t) \neq 0 \\ \mu_0, & \text{if } f(h_t) \leq \bar{k}(\mu_0) \text{ and } s(h_t) = 0 \\ \mu_1, & \text{if } \bar{k}(\mu_0) < f(h_t) \leq \sum_{j=0}^{1} \bar{k}(\mu_j) \text{ and } s(h_t) = 0 \\ \vdots & \vdots \\ \mu_{k-1}, & \text{if } \sum_{i=0}^{k-2} \bar{k}(\mu_i) < f(h_t) \leq \sum_{j=0}^{k-1} \bar{k}(\mu_j) \text{ and } s(h_t) = 0 \\ \mu_k, & \text{if } \sum_{i=0}^{k-1} \bar{k}(\mu_i) < f(h_t) \leq \sum_{j=0}^{k} \bar{k}(\mu_j) \text{ and } s(h_t) = 0 \\ \underline{\mu}, & \text{if } \sum_{i=0}^{k} \bar{k}(\mu_i) < f(h_t) \text{ and } s(h_t) = 0 \end{cases}$$

where $s(h_t) = |\{i \in \{0, 1, ..., t-1\} : y_i = x_H\}|$ and $f(h_t) = |\{i \in \{0, 1, ..., t-1\} : y_i = x_L\}|$.

Notice that this implies two important features of coarse thinkers: (i) the coarse thinker will under-respond to the $\bar{k}(\mu_i)$ first failures observed under μ_i (i.e. every time it observes evidence and doesn't update his belief) and (ii) it will over-respond to the $\bar{k}(\mu_i) + 1^{\text{th}}$ observation (which we will refer to as the pivotal observation) by updating from belief μ_i to belief μ_{i+1}.

We will often compare the belief adopted by the interpretation rule with the posterior that a fine thinker would form upon observing the same evidence. We thus define the function that yields the posterior belief obtained from Bayes updating after observing evidence h_t as:

$$B(h_t) = B(f(h_t), s(h_t)) = \begin{cases} 1, \text{if } s(h_t) \neq 0 \\ \mu_0, \text{if } f(h_t) = s(h_t) = 0 \\ \frac{(1-\pi)^{f(h_t)}\mu_0}{(1-\pi)^{f(h_t)}\mu_0 + 1 - \mu_0}, \text{if } f(h_t) > 0 \text{ and } s(h_t) = 0 \end{cases} \tag{5}$$

2.3. Exploration and Exploitation under Coarse Thinking

The state variable upon which the coarse thinker bases his decision on which action to take at a given period is (μ, n), where $\mu \in \mathfrak{M}$ is the belief used and n is the number of failed attempts observed under this belief. A coarse thinker holding a belief μ that has observed n failures under this belief will then choose either to continue exploring (play R) or begin exploitation (play S).

Coarseness implies a modified dynamic programming problem for the agent, as beliefs are not updated continuously. The values of each state variable for a given belief μ_i will be:

$$\begin{aligned}
&V(\mu_i, 0) = max\{\mu_i\pi(x_H + \delta\frac{E_1}{1-\delta}) + (\mu_i(1-\pi) + (1-\mu_i))(x_L + \delta V(\mu_i, 1)), 0\} \\
&V(\mu_i, 1) = max\{\mu_i\pi(x_H + \delta\frac{E_1}{1-\delta}) + (\mu_i(1-\pi) + (1-\mu_i))(x_L + \delta V(\mu_i, 2)), 0\} \\
&\vdots \\
&V(\mu_i, \bar{k}(\mu_i) - 1) = max\{\mu_i\pi(x_H + \delta\frac{E_1}{1-\delta}) + (\mu_i(1-\pi) + (1-\mu_i))(x_L + \delta V(\mu_i, \bar{k}(\mu_i)), 0\} \\
&V(\mu_i, \bar{k}(\mu_i)) = max\{\mu_i\pi(x_H + \delta\frac{E_1}{1-\delta}) + (\mu_i(1-\pi) + (1-\mu_i))(x_L + \delta V(\mu_{i+1}, 0)), 0\}
\end{aligned} \tag{6}$$

such that, for a given state variable (μ, n), it will play R if $V(\mu, n) > 0$ and play S if $V(\mu, n) = 0$.

Proposition 1. *If $V(\mu_i, 0) > 0$ for a given $\mu_i \in \mathfrak{M}$, a coarse thinker will certainly explore as long as it holds belief μ_i.*

Proof. Imagine value functions as in (6). The expressions for all $V(\mu_i, \cdot)$'s are identical except for the value in the next period conditional on observing a new failure. If $V(\mu_{i+1}, 0)$ is small enough to make $V(\mu_i, \bar{k}(\mu_i)) = 0$ then every other $V(\mu_i, \cdot) = 0$. Conversely, $V(\mu_i, 0) > 0$ only if every other $V(\mu_i, \cdot) > 0$. $\square$

Proposition 1 implies that exploration will only stop upon a change in belief, that is, when the state variable is such as $(\mu, 0)$ for some $\mu \in \mathfrak{M}$.

Proposition 2. *If there is a $\mu_j \in \mathfrak{M}$ for which $V(\mu_j, 0) = 0$ then, for a $\mu_i > \mu_j$, there can only be exploration under μ_i if $\mu_i > p^*$.*

Proof. We know that $p\pi(x_H + \delta\frac{E_1}{1-\delta}) + (p(1-\pi) + (1-p))x_L$ is increasing in p, and we defined p^* as the p for which such expression is equal to 0. As such, we know that if $V(\mu_j, 0) = 0$, $\mu_{j-1}\pi(x_H + \frac{\delta}{1-\delta}E_1) + (\mu_{j-1}(1-\pi) + (1-\mu_{j-1}))x_L$ will be (i) greater than 0 if $\mu_{j-1} > p^*$ or (ii) smaller or equal to 0 if $\mu_{j-1} \leq p^*$. On scenario (ii) we would have $V(\mu_{j-1}, \bar{k}(\mu_{j-1})) = 0$, which by Proposition 1 implies that $V(\mu_{j-1}, 0) = 0$, so that the same logic could be applied for μ_{j-2} and so on. $\square$

Proposition 2 defines $\mu > p^*$ as a necessary condition for exploration if there exists a $\mu' \in \mathfrak{M}$ for which $V(\mu', 0) = 0$.

Proposition 3. *If there is a $\mu_j \in \mathfrak{M}$ greater than p^* for which $V(\mu_j, 0) > 0$ then, for any $\mu_i > \mu_j$, $V(\mu_i, 0) > 0$.*

Proof. For an agent holding state variable $(\mu_{j-1}, \bar{k}(\mu_{j-1}))$, the expected value of exploration is:

$$\mu_{j-1}\pi(x_H + \delta\frac{E_1}{1-\delta}) + (\mu_{j-1}(1-\pi) + (1-\mu_{j-1}))(x_L + \delta V(\mu_j, 0))$$

We know that $V(\mu_j, 0) > 0$, so that whatever the value of μ necessary to make such expression equal to zero is, it must lie in between $(0, p^*)$. As $\mu_{j-1} > \mu_j > p^*$, we know that such expected value of exploration will be positive and that $V(\mu_{j-1}, \cdot) > 0$.

$\square$

Let's denote $\mu_{s-1} = min\{\mu \in \mathfrak{M} : \mu > p*\}$ (or equivalently $\mu_s = max\{\mu \in \mathfrak{M} : \mu \leq p^*\}$). Since we assume that agents can hold belief $\underline{\mu}$ and we know that $V(\underline{\mu}, \cdot) = 0$, by Proposition

2 we know that $V(\mu_j, \cdot) = 0 \forall j \in \{s, s+1, \ldots, k\}$. For μ_{s-1} we know that, since $\mu_{s-1} > p^*$:

$$V(\mu_{s-1}, \bar{k}(\mu_{s-1})) = \mu_{s-1}\pi(x_H + \delta\frac{E_1}{1-\delta}) + (\mu_{s-1}(1-\pi) + (1-\mu_{s-1}))x_L > 0$$

such that by Proposition 1 $V(\mu_{s-1}, 0) > 0$ and then, by Proposition 3, $V(\mu_i, \cdot) > 0 \forall i \in \{0, 1, \ldots, s-1\}$. As such, Propositions 1 to 3 define $\mu > p^*$ as a necessary and sufficient condition for exploration, such that we know that the length of exploration for a coarse thinker will be the amount of time it holds beliefs greater than p^*.

2.4. Stubbornness and Pliancy

Thinking through a coarse space rather than a through a continuum has interesting implications. First, it means that the belief held at each point in time is an imperfect representation of the information the agent possesses. While agents that think in a continuum can map the information they possess to one particular point among infinite other points, coarse thinkers can only map their information to finitely many points, which naturally implies a sense of imperfectness.

Second, it means that agents will not be able to fully extract the information present in the signals they receive. While an agent that thinks in a continuum can finely tune his belief upon observation of new evidence and thus fully extract the information present in it, a coarse thinker that is confronted with new evidence will always face an informational trade-off between remaining inert or doing a coarse and imperfect adjustment to his belief.

As such, coarse thinkers will under-respond to some evidence and over-respond to others. A coarse thinker that has just adopted belief μ_i will remain inert (and thus under-respond) to the $\bar{k}(\mu_i)$ first failures he observed under this belief. Upon observation of the $\bar{k}(\mu_i) + 1^{\text{th}}$ failure he will over-respond by making a coarse adjustment from belief μ_i to belief μ_{i+1}.

High values of $\bar{k}(\mu_i)$ can be then interpreted as stubbornness regarding this belief: the agent requires a lot of evidence that $\theta = 0$ in order to lower his expectations about θ being 1. Low values of $\bar{k}(\mu_i)$, on the other hand, are associated with pliancy: the agent requires just a few evidence to abandon his current belief in favor of a smaller one. Different interpretation rules will define different sequences $(\bar{k}(\mu_i))_{i=0}^{k}$ and thus different balances between stubbornness and pliancy for each belief in his belief set, and such balance will determine the way the learning process will develop.

As we have seen in the former section, coarseness does not imply a difference in optimal behavior given belief held compared to agents that think in a continuum: both coarse and fine thinkers explore when they hold beliefs greater than p^* and exploit when they hold beliefs smaller or equal to p^*. What is different in their behavior, though, is the way they interpret signals. In the next section we will further discuss signal interpretation for coarse thinkers and analyze, for particular classes of interpretation rules, what are their implications on the learning process of agents and on how coarse thinkers would behave under a exploration-exploitation dilemma.

3. Interpretation Rules

Interpretation is treated in this paper in a very general but intuitive way. Just like the sailor that adopts a belief about whether or not he is approaching land as a function of his observation on the pattern of the swells in the sea or the behavior of other animals around him, the coarse thinker will adopt a belief about the state of nature θ as a function of the evidence h_t available to him at that time.

This generality allows for many biases. A superstitious sailor that has dreamed about reaching an island filled with fruits and fresh water might be very confident about approaching land and ignore all evidence of the opposite, just like a stubborn coarse thinker with an arbitrarily large $\bar{k}(\mu_0)$ will effectively ignore all evidence of θ in fact being zero.

A high degree of stubbornness implies the constant underestimation of the informational content present in the signals the coarse thinker observes, and thus causes the belief adopted by it to be systematically above the belief of the fine thinker. That is, for any history h_t yielding $B(h_t) \in [\mu_{i+1}, \mu_i)$, the interpretation would always result in $\mu(h_t) = \mu_j$ with $j \in \{0, 1, \ldots, i\}$.

Conversely, an overly pliant interpretation rule would result in the overestimation of the informational content present in the signals, such that the belief of the coarse thinker would be systematically below the belief of the fine thinker. That is, for a history h_t yielding $B(h_t) \in (\mu_{i+1}, \mu_i]$, the interpretation would give $\mu(h_t) = \mu_j$ with $j \in \{i+1, i+2, \ldots, k\}$.

Even though there are infinite possible interpretation rules, including those that always under or over-estimate the informational content present in signals, there is a class of interpretation rules that carries a sense of centrality. Imagine a history h_t such that $B(h_t) = \mu_i$ and another history h'_t such that $B(h'_t) \in (\mu_{i+1}, \mu_i)$. The "central" interpretation rules are those that would assign $\mu(h_t) = \mu_i$ and $\mu(h'_t) \in \{\mu_i, \mu_{i+1}\}$.

Notice that these interpretation rules do not incur in a systematic error other than the natural misrepresentation that arises from coarseness: the information available in history h'_t is misrepresented by the adoption of belief μ_i or μ_{i+1}, but whenever the fine thinker reaches a belief belonging to $\mathfrak{M}$ the coarse thinker will be holding that belief as well. The belief decay implied by these interpretation rules over time will not be detached from the belief decay curve of fine thinkers, but will rather zigzag it, such that the coarse belief is sometimes higher and sometimes lower than the fine belief.

Suppose that a fine thinker reaches belief $\mu_i \in \mathfrak{M}$ at a time $t = m$. A coarse thinker with an interpretation rule of this type will always shift from belief μ_{i-1} to belief μ_i at a time $t \leq m$, and will always shift away from belief μ_i to belief μ_{i+1} at a time $t > m$. The fact that the coarse belief decay roughly follows the bayesian decay allows these interpretation rules to be interpreted in categorical terms: they will behave just like a fine thinker whose belief space $[0, 1]$ is partitioned into categories, such that every posterior under a given category is treated as the same belief. This class of interpretation rules will be analyzed in section 3.1, while the rules that do incur in a systematic error will be treated in section 3.2.

The key aspect we are interested in studying is the extent of over or under exploration that is implied by an interpretation I. As we have seen, each interpretation rule will define a sequence $(\bar{k}(\mu_i))_{i=0}^{k}$. Consider $\mu_s = max\{\mu \in \mathfrak{M} : \mu \leq p^*\}$, the highest belief in $\mathfrak{M}$ that is lower or equal to p^*. The length of exploration under interpretation rule I (i.e. the amount of failures the coarse thinker will observe before shifting to exploitation) will then be $L(I) = \sum_{i=0}^{s-1} \bar{k}(\mu_i)$. We can thus define the amount of over/under-exploration in a given interpretation rule I to be:

$$A(I) = \sum_{i=0}^{s-1} \bar{k}(\mu_i) + 1 - B^{-1}(p^*) \tag{7}$$

such that $A(I) > 0$ when the coarse thinker over-explores and $A(I) < 0$ when the coarse thinker under-explores.

Notice that, given the fact that interpretation rules pick a sequence $(\bar{k}(\mu_i))_{i=0}^{k}$ belonging to the set of k-dimensional non-negative integers, which is infinite but bounded from below, there exists a lower bound (but not an upper bound) on $A(I)$. Consider the most pliant of interpretation rules, the one that leads to a shift to the next belief every time the agent observes a failure. He will then hold beliefs larger than p^* for the $s-1$ first failures he observes, and will stop exploration upon observing the s^{th} failure.

For agents with a given belief set $\mathfrak{M}$, the amount of under-exploration under such interpretation rule will correspond to a bound on under-exploration:

$$\underline{A} = s - B^{-1}(p^*) \tag{8}$$

In order to simplify notation, in the remaining of this section I will note $B(f(h_t) = i, s(h_t) = 0)$ as $B(i)$.

3.1. Categorical Interpretation Rules

Coarseness has been previously modelled in theoretical work (Mullainathan [21], Mullainathan, Schwartzstein and Schleifer [11]) as the product of a categorization on belief space $[0, 1]$. Although it is true that the concept of categorization naturally implies imposing a coarser granularity to a set, coarse thinking does not need to presuppose the categorization of an underlying space. In fact, such supposition already relies on assumptions on how the agents interpret information.

As we have previously seen, any interpretation rule I will define a partition on the space of possible evidence observed at a given period $\{x_L, x_H\}^t$. Let's restrict our attention to the sub-space of $\{x_L, x_H\}^t$ in which no success was realized (i.e. the space of realized payoffs that could possibly lead to exploration in our problem). The function $B(f(h_t))$ yielding the Bayesian posterior after observing a history h_t with $f(h_t)$ failures and no successes defines a bijection between this sub-space and $[0,1]$, such that, under two conditions, we can understand the interpretation rules as defining a convex partition on the posterior space of an implied bayesian thinker.

Let's imagine a categorical bayesian: an agent whose posterior space is partitioned into categories, with each category being associated with a specific belief. Formally, this would account for an agent with category space C and categorization function $c : [0,1] \to C$, where each category in C adopts a belief q_c belonging to the set of posteriors that are mapped into that category.

A coarse thinker that shifts from belief μ_0 to belief μ_1 after observing the $\bar{k}(\mu_0) + 1^{\text{th}}$ failure will do so at the same time as a categorical bayesian with $B(\bar{k}(\mu_0)) = sup\{p \in C_1\}$. It will then shift from belief μ_1 to μ_2 after observing $\bar{k}(\mu_1)$ further failures, just like the categorical bayesian with $B(\bar{k}(\mu_0) + \bar{k}(\mu_1)) = sup\{p \in C_2\}$.

A coarse thinker with an interpretation rule I defining a given sequence $(\bar{k}(\mu_i))_{i=0}^{k}$ will then behave identically as a categorical bayesian with posterior space partitioned as

$$\begin{aligned}
C_0 &= [B(\bar{k}(\mu_0)), 1) \\
C_i &= [B(\sum_{j=0}^{i} \bar{k}(\mu_j)), B(\sum_{j=0}^{i-1} \bar{k}(\mu_j))) \quad \forall i \in \{1, 2, ..., k\} \\
\underline{C} &= [0, B(\sum_{j=0}^{k} \bar{k}(\mu_j))]
\end{aligned} \tag{9}$$

As long as $\mu_i = q_{c_i} \in C_i \forall i \in \{1, 2, \ldots, k\}$, i.e. as long as the belief associated to each category belongs to that category.

It is easy to see that, in order for belief μ_i to belong to category C_i, μ_i must lie in between $sup\{p \in C_i\}$ and $max\{p \in C_{i+1}\}$. As such, any interpretation rule defining a sequence $(\bar{k}(\mu_i))_{i=0}^{k}$ will allow for a categorical representation with categories as defined in (9) if it satisfies:

$$\sum_{j=0}^{i-1} \bar{k}(\mu_j) + 1 \leq B^{-1}(\mu_i) \quad \forall i \in \{1, 2, \ldots, k\} \tag{10}$$

$$\sum_{j=0}^{i} \bar{k}(\mu_j) + 1 > B^{-1}(\mu_i) \quad \forall i \in \{1, 2, \ldots, k\} \tag{11}$$

Condition (10) defines a bound on the stubbornness implied by the interpretation rule, stating that the coarse thinker should not be overly stubborn so as to still be holding belief μ_{i-1} at the time the fine thinker adopts belief μ_i; whereas condition (11) defines a bound on the pliancy implied by the interpretation rule, stating that the coarse thinker should not be overly pliant so as to have already abandoned belief μ_i at the time the fine thinker adopts it.

Proposition 4. *Any interpretation rule that allows for a categorical representation will induce $A(I) > 0$ if p^* is categorized under C_{s-1} and $A(I) \leq 0$ if p^* is categorized under C_s.*

Proof. Categorizing p^* under C_{s-1} implies that $\sum_{i=0}^{s-1} \bar{k}(\mu_i) + 1 \in (B^{-1}(p^*), B^{-1}(\mu_s)]$, which in turn implies that $A(I) \in (0, B^{-1}(\mu_s) - B^{-1}(p^*)]$.

Conversely, categorizing p^* under C_s implies that $\sum_{i=0}^{s-1} \bar{k}(\mu_i) + 1 \in (B^{-1}(\mu_{s-1}), B^{-1}(p^*)]$, which in turn implies that $A(I) \in (B^{-1}(\mu_{s-1}) - B^{-1}(p^*), 0]$. $\square$

By Proposition 4 we know that any interpretation rule that allows for a categorical representation will lead to over-exploration when p^* is categorized under C_{s-1} and under or just-exploration when p^* is categorized under C_s. This is intuitive: if the agent understands the cut-off belief p^* and its neighbours as "just sufficiently likely", then it will certainly be still exploring once it reaches p^* and thus the error in exploration will certainly be positive. On the other hand, if the agent understands the cut-off belief p^* as "likely but just not enough", then it will not be exploring anymore once it reaches p^* and thus the error in exploration will almost surely be negative (it can also be 0 if p^* is the highest belief in that category).

The bounds on potential error in exploration can give us an idea on how coarseness relates to the extent of over or under-exploration under such interpretation rules.

Proposition 5. *The error in exploration resulting from any interpretation rule that allows for a categorical representation can be made arbitrarily low as the partition of the implied belief space $[0, 1]$ gets finer around p^*.*

Proof. The maximum error in interpretation for such interpretation rules is $B^{-1}(\mu_s) - B^{-1}(p^*)$ when p^* is categorized under C_{s-1} and $B^{-1}(\mu_{s-1}) - B^{-1}(p^*)$ when p^* is categorized under C_s. As the categorization gets finer around p^* we will have:

$$\lim_{\mu_{s-1} \to p^{*-}} B^{-1}(\mu_{s-1}) - B^{-1}(p^*) = \lim_{\mu_s \to p^{*+}} B^{-1}(\mu_s) - B^{-1}(p^*) = 0$$

$\square$

3.1.1. Optimal Interpretation Rule

We can understand the optimal interpretation rule as the one that maps every possible history h_t into the belief that the fine thinker would pick if she was constrained to choose only between the beliefs belonging to $\mathfrak{M}$.

It is easy to see that such interpretation rule has to satisfy the conditions for categorical representation defined in the earlier section. We know that for a history h_t yielding $B(h_t) = \mu_i$, a fine thinker would choose $\mu(h_t) = \mu_i$ and that for any history $h_{t'}$ yielding a $B(h_{t'}) \in (\mu_i, \mu_{i+1})$, a fine thinker would choose either μ_i or μ_{i+1}. We thus know that the optimal interpretation rule will define for every $i \in \{0, 1, ..., k\}$ a $\bar{k}(\mu_i)$ such that $\sum_{j=0}^{i} \bar{k}(\mu_j) + 1 \in (B^{-1}(\mu_i), B^{-1}(\mu_{i+1})]$, which is equivalent to conditions (10) and (11).

We then know by Proposition 4 that the optimal interpretation rule will lead to over-exploration if the partition of $[0, 1]$ defined by it is such that p^* is categorized along with μ_{s-1}, or to under-exploration in the case the cut-off belief is categorized along with μ_s. We can, however, narrow these results down a bit.

We know that upon observing history h_t, a coarse thinker commits an interpretation error $e(h_t) = \mu(h_t) - B(h_t)$. We can define the degree of misinterpretation during the explorative phase implied in a given interpretation rule as:

$$D(I) = \sum_{i=0}^{L(I)} \mu(f(h_t) = i, s(h_t) = 0) - B(f(h_t) = i, s(h_t) = 0) \tag{12}$$

Which sums for every period the distance between the belief adopted by the interpretation rule and the belief that a fine thinker would form if confronted with the same evidence.

The optimal interpretation rule can then be defined as:

$$(\bar{k}^*(\mu_i))_{i=0}^k \in \underset{(\bar{k}(\mu_i))_{i=0}^k}{\arg\min} |D(I)| \tag{13}$$

subject to constraints (10) and (11).

This interpretation rule is equivalent to choosing, for every history h_t with $f(h_t)$ failures and 0 successes, the belief in $\mathfrak{M}$ that is closer to $B(f(h_t))$.

Considering $\bar{\mu}_{i,i+1} = \frac{\mu_i + \mu_{i+1}}{2}$, the optimal interpretation rule will thus define $(\bar{k}(\mu_i))_{i=0}^k$ such that, for every $i \in \{0, 1, ..., k\}$, $\sum_{j=0}^{i} \bar{k}(\mu_j) + 1 = B^{-1}(\bar{\mu}_{i,i+1})$. More specifically, this interpretation rule implies that agents stop exploring when $f(h_t) = B^{-1}(\bar{\mu}_{s-1,s})$.

The amount of over/under-exploration will be then given by:

$$A(I) = B^{-1}(\bar{\mu}_{s-1,s}) - B^{-1}(p^*) \tag{14}$$

As such, over (under) exploration will happen when p^* is higher (lower) than $\bar{\mu}_{s-1,s}$. Notice that, since the optimal interpretation rule satisfies the conditions for categorical representation and defines $\bar{\mu}_{s-1,s} = max\{p \in C_s\}$, p^* being higher (lower) than $\bar{\mu}_{s-1,s}$ is equivalent to it being categorized under C_{s-1} (C_s).

If the cut-off belief belongs to the set of coarse beliefs that the agent can hold (i.e. if $\mu_s = p^* \in \mathfrak{M}$), the optimal interpretation will imply $A(I) \leq 0$. Whether p^* belongs to $\mathfrak{M}$ or not, by Proposition 5 we know that over/under-exploration gets arbitrarily close to 0 as the partition of the posterior space gets finer around p^*.

3.1.2. Punctual Interpretation Rule

Imagine an interpretation rule that always adopts belief μ_i at the same time as the bayesian thinker would adopt the same belief, such that $\sum_{j=0}^{i-1} \bar{k}(\mu_j) + 1 = B^{-1}(\mu_i) \forall i \in \{1, 2, \ldots, k-1\}$. This interpretation rule is the most stubborn among the class of categorical interpretation rules: the belief adopted by it is always either the same or above the belief of the fine thinker.

The belief decay implied in this rule is not detached from the bayesian belief decay, since both curves meet periodically (whenever the fine thinker reaches a belief $\mu \in \mathfrak{M}$). But since $\mu(h_t) \geq B(h_t)$ for any possible history, its degree of misinterpretation $D(I)$ will be persistently positive.

The categorization implied by this interpretation rule is one in which $\mu_i = max\{p \in C_i\} \forall i \in \{1, 2, \ldots, k\}$. In particular, $\mu_s = max\{p \in C_s\}$, such that we know that p^* will be categorized under C_{s-1} whenever it does not belong to $\mathfrak{M}$. That is, when $p^* \notin \mathfrak{M}$ this interpretation rule will certainly lead to over-exploration of magnitude $A(I) = B^{-1}(\mu_s) - B^{-1}(p^*) > 0$, whereas when $p^* \in \mathfrak{M}$ there will be just-exploration.

3.2. Detached Interpretation Rules

In the last section we analyzed interpretation rules that roughly followed the decay of a bayesian belief. We have seen that such interpretation rules are bounded on their stubbornness and pliancy, which prevents the resulting beliefs from becoming too detached to the belief that a fine thinker would form if confronted with the same evidence. We now look at interpretation rules that systematically violate conditions (10) and (11) and that thus lead to a belief decay that is detached from the bayesian one.

3.2.1. Over-Stubbornness/Anchoring

An interpretation rule that systematically violates condition (10) will always underestimate the informational content of the signals it observes, such that the resulting belief will always be above that of the fine thinker. Violating the stubbornness condition for every belief would account to defining a sequence $(\bar{k}(\mu_i))_{i=0}^{k}$ such that:

$$\sum_{j=0}^{i-1} \bar{k}(\mu_j) + 1 > B^{-1}(\mu_i) \forall i \in \{1, 2, \ldots, k\} \tag{15}$$

That is, for every history h_t with $B(h_t) \in [\mu_{i+1}, \mu_i)$, the anchored interpretation rule would yield a $\mu(h_t) = \mu_j$ with $j \in \{0, 1, \ldots, i\}$. Such systematic bias can be thought of in terms of the anchoring bias discussed in the behavioral literature (Tversky and Kahneman [19], Chapman and Johnson [22]). Beliefs that are anchored around a certain value will exhibit a greater inertia and insufficient adjustments to later signals. The sailor that has dreamt about reaching a tropical island, for example, might anchor his beliefs about its existence and underestimate the value of signals that indicate otherwise.

It is thus intuitive that such interpretation would lead him to over-explore the unknown sea in his search. More specifically, we know that agents with such stubbornness will always reach a belief μ_i with $i \in \{1, 2, \ldots, k\}$ after the fine thinker has reached μ_i, such that the extent of over exploration would lie in:

$$A(I) \in (B^{-1}(\mu_s) - B^{-1}(p^*), \infty)$$

3.2.2. Over-Pliancy

An interpretation rule that systematically violates condition (11) would always over-estimate the informational value of the evidence it encounters, such that the belief adopted by it would be always below that of a fine thinker. An interpretation rule that violates the pliancy condition would then define a sequence $(\bar{k}(\mu_i))_{i=0}^{k}$ such that:

$$\sum_{j=0}^{i} \bar{k}(\mu_j) + 1 \leq B^{-1}(\mu_i) \forall i \in \{1, 2, \ldots, k\} \tag{16}$$

For any history h_t yielding a $B(h_t) \in (\mu_{i+1}, \mu_i]$, this interpretation rule would adopt a $\mu(h_t) \in \{\mu_{i+1}, \mu_{i+2}, \ldots, \mu_k, \underline{\mu}\}$. Since the pliant coarse thinker will always abandon a given belief belonging to $\mathfrak{M}$ before the fine thinker adopts it, it is clear that such interpretation rules will always exhibit under-exploration.

As we have already previously mentioned, pliancy is bounded for coarse thinkers. The most pliant of interpretation rules is the one that shifts to the next belief available every time it observes a failure. The under-exploration implied by such interpretation rule is then $A(I) = \underline{A} = s - B^{-1}(p^*)$.

As such, pliant interpretation rules would always imply under-exploration in the interval:

$$A(I) \in [s - B^{-1}(p^*), B^{-1}(\mu_{s-1}) - B^{-1}(p^*)]$$

We can observe that the minimum extent of under-exploration in this case coincides with the maximum extent of under-exploration possible under a categorical interpretation rule. Similarly, the minimum extent of over-exploration under a stubborn interpretation rule will coincide with the maximum extent of over-exploration possible in a categorical interpretation. These results can be summarized as follows:

$$\begin{aligned} A(I^P) < 0 \text{ and } A(I^P) &\in [s - B^{-1}(p^*), B^{-1}(\mu_{s-1}) - B^{-1}(p^*)] \\ A(I^{C_s}) \leq 0 \text{ and } A(I^{C_s}) &\in (B^{-1}(\mu_{s-1}) - B^{-1}(p^*), 0] \\ A(I^{C_{s-1}}) > 0 \text{ and } A(I^{C_{s-1}}) &\in (0, B^{-1}(\mu_s) - B^{-1}(p^*)] \\ A(I^S) > 0 \text{ and } A(I^S) &\in (B^{-1}(\mu_s) - B^{-1}(p^*), \infty) \end{aligned} \tag{17}$$

4. Concluding Remarks

Modeling coarseness of beliefs poses a challenge in terms of how to represent evidence interpretation. For fine thinkers, such process is smooth: they can always finely adjust their beliefs based on the information contained in the signals they encounter. Coarse thinkers, however, are restricted to representing their information through finitely-many beliefs. Furthermore, whenever they are exposed to new evidence they will face a trade-off between maintaining their current belief or performing a coarse and imperfect adjustment. As such, each interpretation rule will define a particular balance between stubbornness and pliancy regarding each belief.

Although a coarse thinker would adopt the same strategy as a fine thinker in a two-armed bandit problem, the fact that he conditions his actions on an imperfect measure of his information should affect the way he allocates resources into exploration. We analyzed such impact for different types of interpretation: agents that underestimate the informational value of the evidence they observe would incur in over-exploration, whereas agents that overestimate the information contained in their evidence would under-explore. These results highlight the impact that imbalances in the stubbornness and pliancy implied in an interpretation rule might have in the learning process of a coarse thinker.

For agents whose interpretation follows roughly that of a fine thinker, the extent of exploration will depend on how they perceive the cut-off belief. As we have seen, such interpretations imply that the coarse thinker can be treated as an agent whose continuum belief space is partitioned into coarse categories, such that every posterior belonging to a given category is perceived equally. As such, exploration length would depend on whether the cut-off belief is categorized into a belief that still treats the good state of nature as sufficiently likely or not. This result highlights the impact that the imperfect representation of information through coarse beliefs might have in an exploration setting.

References

[1] A. V. Banerjee, A simple model of herd behavior, The Quarterly Journal of Economics 107 (3) (1992) 797–817.
URL `http://dx.doi.org/10.2307/2118364`

[2] S. Bikhchandani, D. Hirshleifer, I. Welch, A theory of fads, fashion, custom, and cultural change as informational cascades, Journal of Political Economy 100 (5) (1992) 992–1026.
URL `http://www.jstor.org/stable/2138632`

[3] D. Acemoglu, M. A. Dahleh, I. Lobel, A. Ozdaglar, Bayesian learning in social networks, The Review of Economic Studies 78 (4) (2011) 1201–1236. doi:10.1093/restud/rdr004.
URL `http://dx.doi.org/10.1093/restud/rdr004`

[4] V. P. Crawford, J. Sobel, Strategic information transmission, Econometrica 50 (6) (1982) 1431–1451.
URL `http://www.jstor.org/stable/1913390`

[5] J. Farrell, M. Rabin, Cheap talk, Journal of Economic Perspectives 10 (3) (1996) 103–118. doi:10.1257/jep.10.3.103.
URL `http://www.aeaweb.org/articles?id=10.1257/jep.10.3.103`

[6] E. Kamenica, M. Gentzkow, Bayesian persuasion, American Economic Review 101 (6) (2011) 2590–2615. doi:10.1257/aer.101.6.2590.
URL `http://www.aeaweb.org/articles?id=10.1257/aer.101.6.2590`

[7] D. Bergemann, S. Morris, Information design, bayesian persuasion, and bayes correlated equilibrium, American Economic Review 106 (5) (2016) 586–91. doi:10.1257/aer.p20161046.
URL `http://www.aeaweb.org/articles?id=10.1257/aer.p20161046`

[8] P. Bolton, C. Harris, Strategic experimentation, Econometrica 67 (2) (1999) 349–374. doi:10.1111/1468-0262.00022.
URL `https://onlinelibrary.wiley.com/doi/abs/10.1111/1468-0262.00022`

[9] G. Keller, S. Rady, M. Cripps, Strategic experimentation with exponential bandits, Econometrica 73 (1) (2005) 39–68. doi:10.1111/j.1468-0262.2005.00564.x.
URL `https://onlinelibrary.wiley.com/doi/abs/10.1111/j.1468-0262.2005.00564.x`

[10] J. Roland G. Fryer, P. Harms, Two-armed restless bandits with imperfect information: Stochastic control and indexability, Working Paper 19043, National Bureau of Economic Research (May 2013). doi:10.3386/w19043.
URL `http://www.nber.org/papers/w19043`

[11] S. Mullainathan, J. Schwartzstein, A. Shleifer, Coarse thinking and persuasion, The Quarterly Journal of Economics 123 (2) (2008) 577–619. doi:10.1162/qjec.2008.123.2.577.
URL `http://dx.doi.org/10.1162/qjec.2008.123.2.577`

[12] A. Guarino, P. Jehiel, Social learning with coarse inference, American Economic Journal: Microeconomics 5 (1) (2013) 147–74. doi:10.1257/mic.5.1.147.
URL `http://www.aeaweb.org/articles?id=10.1257/mic.5.1.147`

[13] J. A. Bohren, D. Hauser, Bounded rationality and learning: A framework and a robustness result, PIER Working Paper No. 17-007.
URL `https://ssrn.com/abstract=2968460`

[14] C. N. Macrae, G. V. Bodenhausen, Social cognition: Thinking categorically about others, Annual Review of Psychology 51 (1) (2000) 93–120, pMID: 10751966. doi:10.1146/annurev.psych.51.1.93.
URL `https://doi.org/10.1146/annurev.psych.51.1.93`

[15] A. B. Markman, D. Gentner, Thinking, Annual Review of Psychology 52 (1) (2001) 223–247, pMID: 11148305. doi:10.1146/annurev.psych.52.1.223.
URL `https://doi.org/10.1146/annurev.psych.52.1.223`

[16] G. Murphy, B. Ross, Predictions from uncertain categorizations, Cognitive Psychology 27 (2) (1994) 148 – 193. doi:https://doi.org/10.1006/cogp.1994.1015.
URL `http://www.sciencedirect.com/science/article/pii/S0010028584710152`

[17] B. Malt, G. Murphy, B. Ross, Predicting features for members of natural categories when categorization is uncertain, Journal of Experimental Psychology: Learning, Memory, and Cognition 21 (1995) 646–661.

[18] D. Kahneman, A. Tversky, Subjective probability: A judgment of representativeness, Cognitive Psychology 3 (3) (1972) 430 – 454. doi:https://doi.org/10.1016/0010-0285(72)90016-3.
URL `http://www.sciencedirect.com/science/article/pii/0010028572900163`

[19] A. Tversky, D. Kahneman, Judgment under uncertainty: Heuristics and biases, Science 185 (4157) (1974) 1124–1131. doi:10.1126/science.185.4157.1124.
URL `http://science.sciencemag.org/content/185/4157/1124`

[20] J. Krueger, R. W. Clement, Memory-based judgments about multiple categories: A revision and extension of tajfel's accentuation theory., Journal of Personality and Social Psychology 67(1) (1994) 35–47.

[21] S. Mullainathan, Thinking through categories, mimeo, MIT.

[22] G. Chapman, E. Johnson, Incorporating the irrelevant: Anchors in judgments of belief and value, In T. Gilovich, D. Griffin, D. Kahneman (Eds.), Heuristics and biases (2002) 120–138.