Mental Processes and Decision Making*

Olivier Compte
PSE, Paris

Andrew Postlewaite
University of Pennsylvania

March 2009

PRELIMINARY AND INCOMPLETE; PLEASE DO NOT QUOTE

*Compte: Paris School of Economics, 48 Bd Jourdan, 75014 Paris (e-mail: compte@enpc.fr); Postlewaite: Department of Economics, 3718 Locust Walk, University of Pennsylvania, Philadelphia, PA 19104-6207 (e-mail: apostlew@econ.upenn.edu). A large part of this work was done while Postlewaite visited the Paris School of Economics; their hospitality is gratefully acknowledged. We thank David Dillenberger, Kfir Eliaz, Yuval Salant and Alvaro Sandroni for helpful discussions. The authors thank the Gould Foundation for financial support, and Postlewaite thanks the National Science Foundation for financial support.

Abstract

It ain't so much the things we don't know that get us into trouble. It's the things that we know that just ain't so.
Artemus Ward

1 Introduction

Robert is convinced that he has ESP and offers the following statement to support this belief: "I was thinking of my mother last week and she called right after that." Robert is not alone in his beliefs; more people believe in ESP than in evolution, and in the U.S. there are twenty times as many astrologers as there are astronomers.[1] Readers who don't believe in ESP might dismiss Robert and other believers as undereducated anomalies, but there are sufficiently many other similar examples to give pause. Nurses who work on maternity wards believe (incorrectly) that more babies are born when the moon is full[2], and it is widely believed that infertile couples who adopt a child are subsequently more likely to conceive than similar couples who did not adopt (again, incorrectly).[3]

We might simply decide that people that hold such beliefs are stupid or gullible, at the risk of finding ourselves so described for some of our own beliefs.[4] Whether or not we are so inclined, many economic models have at their core a decision-making module, and those models must somehow take account of agents' beliefs, however unsound we may think them.

Our interest in the widespread belief in ESP goes beyond the instrumental concern for constructing accurate decision making modules for our models. The deeper question is why people hold such questionable beliefs? The simple (simplistic?) response that a large number of people are stupid is difficult to accept given the powerful intellectual tools that evolution has provided us in many domains. How is that evolution has generated a brain that can scan the symbols on a page of paper and determine which subway connected to which bus will systematically get someone to work on time, and yet believe in ESP?

Our aim in this paper is to reconcile the systematic mistakes we observe in the inferences people draw from their experiences with evolutionary forces that systematically reward good decisions. We will lay out a model of how an individual processes streams of informative signals that is (a) optimal, and (b) leads to incorrect beliefs such as Robert's. The reconciliation is possible because of computational bounds we place on mental processing. Roughly speaking, our restrictions on mental processing preclude an agent from recalling every signal he receives perfectly. Consequently, he must employ some sort of summary

[1] See Gilovich (1991), page 2.
[2] See G. O. Abell and B. Greenspan (1979).
[3] See E. J. Lamb and S. Leurgans (1979).
[4] There are numerous examples of similarly biased beliefs people hold. Research has demonstrated that people frequently estimate the connection between two events such as cloud seeding and rain mainly by the number of positive-confirming events, that is, where cloud seeding is followed by rain. Cases of cloud seeding and no rain and rain without cloud seeding tend to be ignored (Jenkins and Ward (1965) and Ward and Jenkins (1965).)

statistics that capture as well as possible the information content of all the signals that he as seen. We assume that agents do not have distinct mental processes for every problem the might face, hence an optimal process will do well for "typical" problems, but less well for "unusual" problems. Given the restrictions agents face in our model, they optimally ignore signals that are very uninformative. Robert's experience of his mother calling right after he thought of her is quite strong evidence in support of his theory that he has ESP. His problem lies in his having not taken into account the number of times his mother called when he *hadn't* thought of her. Such an event may have moved Robert's posterior belief that he had ESP only slightly, but the accumulation of such small adjustments would likely have overwhelmed the small number of instances which seem important. Our primary point is that the mental processing property that we suggest leads Robert to conclude that he has ESP – ignoring signals that by themselves have little information – will, in fact, be optimal when "designing" a mental process that must be applied to large sets of problems when there are computational bounds.

We lay out our model of mental processes in the next section. The basic model is essentially that analyzed by Wilson (2004) and by Cover and Hellman (1970), but our interest differs from that of those authors. In those papers, as in ours, it is assumed that an agent has bounded memory, captured by a set of mental states. Agents receive a sequence of signals that are informative about which of two possible states of Nature is the true state. The question posed in those papers is how the agent can optimally use the signals to move among the finite set of mental states, knowing that at some random time he will be called upon to make a decision, and his current mental state is all the information he has about the signals he has received.

Cover and Hellman and Wilson characterize the optimal way to transit among mental states as additional signals arrive when the expected number of signals the agent receives before making a decision goes to infinity. Our interest differs from these authors in two respects. First, as mentioned above, our point of view is that an agent's mental system – the set of mental states and the transition function – have evolved to be optimal for a class of problems rather than being designed for a single specific problem. Second, we are interested in the case that the expected number of signals that an agent will receive before making a decision is bounded. We compare different mental systems in section 3 and discuss robustness of our analysis in section 4.

2 The model

Decision problem. There are two states, $\theta = 1, 2$. The true state is $\theta = 1$ with probability π_0. An agent receives a sequence of signals imperfectly correlated with the true state, that he will use to take a single decision. The decision is a choice between two alternatives, $a \in \{1, 2\}$. To fix ideas, we assume the following payoff matrix, where $g(a, \theta)$ is the payoff to the agent when he takes action a in state θ:

$g(a,\theta)$	1	2
1	1	0
2	0	1

There are costs c_1 and c_2 associated with decisions 1 and 2 respectively. Let $c = (c_1, c_2)$ denote the cost profile, and $u(a, \theta, c)$ the utility associated with each decision a when the state is θ and cost profile is c. We assume that the utility function takes the form

$$u(a, \theta, c) = g(a, \theta) - c_a.$$

The cost c is assumed to be drawn from a distribution with full support on $[0,1] \times [0,1]$. The cost vector c is known to the agent prior to the decision. It is optimal to choose $a = 1$ when the agent's belief $\pi \geq \frac{1+c_1-c_2}{2}$. In what follows, we let $v(\pi)$ denote the payoff the agent derives from optimal decision making when π is his belief that the true state is $\theta = 1$. We have:

$$v(\pi) = E_{c_1,c_2} \max\{\pi - c_1, 1 - \pi - c_2\}.$$

It is straightforward to show that v is strictly convex. For example, if costs are uniformly distributed, $v(\pi) = v(1-\pi)$ and a calculation shows that for $\pi \geq 1/2$, $v(\pi) = \frac{1}{6} + \frac{4}{3}(\pi - \frac{1}{2})^2(2 - \pi)$ (see figure 1).

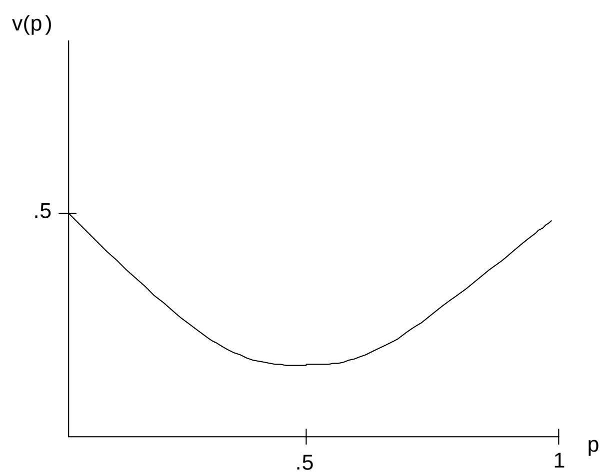

Figure 1: $v(\pi)$ when costs are distributed uniformly

The signals received. Signals are drawn independently, conditional on the true state θ, from the same distribution with continuous density $f(\cdot \mid \theta)$. It will be convenient to denote by $l(x)$ the likelihood ratio defined by:

$$l(x) = \frac{f(x \mid \theta = 1)}{f(x \mid \theta = 2)}.$$

We assume that signals are received over time, at dates $t = 0, 1, ...$, and that the decision must be taken at some random date $\tau \geq 1$. For simplicity, we assume that τ follows an exponential distribution with parameter $1 - \lambda$:

$$P(\tau = t \mid \tau \geq t) = 1 - \lambda.$$

This assumption is meant to capture the idea that the agent will have received a random number of signals prior to making his decision. The fact that this number is drawn from an exponential distribution is not important, but makes computations tractable. The parameter λ provides a measure of the number of signals the agent is likely to receive before he must make a decision: the closer λ is to 1, the larger the expected number of signals. Note that the agent always receives at least one signal.

Limited perceptions. We do not take for granted that players correctly perceive likelihood ratios. For the most part of the paper, we will assume that they do, but we wish to keep in mind that they might not be able to correctly perceive likelihood ratios. We shall examine in Section 4 how our approach extends to that case. Formally, we denote by $\widetilde{l}$ the perceived likelihood ratio, and for now assume that

$$\widetilde{l}(x) = l(x) \text{ for all } x.$$

Below, we shall incorporate noisy perceptions of the following form:

$$\widetilde{l}(x) - 1 = \widetilde{\mu}(l(x) - 1) \text{ when } l(x) > 1$$

and

$$\frac{1}{\widetilde{l}(x)} - 1 = \widetilde{\mu}(\frac{1}{l(x)} - 1) \text{ when } l(x) < 1,$$

where $\widetilde{\mu}$ is a random variable taking values in $(0, +\infty)$. The case in which $E\widetilde{\mu} = 1$ corresponds to the case in which an agent has noisy but unbiased perceptions of the ratio $(\frac{1}{l} - 1)$, conditional on x. The interpretation is that agents correctly perceive whether x is evidence for $\theta = 1$ (i.e. $l(x) > 1$), or against ($l(x) < 1$), but they have limited ability to assess the strength of evidence. For example, the degree to which the result of a medical test is evidence for a particular illness might depend on many factors, and the practitioner may pay attention to a limited number of them.

Limited information processing. We do not take for granted, however, that agents can finely record and process information. Agents are assumed to have a limited number of *states of mind*, and each signal the agent receives is assumed to (possibly) trigger a change in his state of mind. We have in mind, however, that those transitions apply across many decision problems that the agent may face, so the transition will not be overly problem-specific or tailored to the particular decision problem at hand. Thus, we shall assume that transitions may only depend on the perceived likelihood ratio associated with the signal received. Formally a state of mind is denoted $s \in S$, where S is a finite set.

For any signal x received, changes in state of mind depend on the perceived likelihood ratio $\widetilde{l}$ associated with x. We denote by T the transition function:

$$s' = T(s, \widetilde{l}).$$

To fix ideas, we provide a simple example. We will later generalize the approach.

Example 1: The agent may be in one of three states of mind $\{s_0, s_1, s_{-1}\}$. His initial state is s_0. When he receives a signal x, he compares the likelihood ratio $\widetilde{l}$ with 1. The event $A_0^+ = \{\widetilde{l} > 1\}$ is interpreted as evidence in favor of state $\theta = 1$, while $A_0^- = \{\widetilde{l} < 1\}$ is interpreted as evidence against $\theta = 1$. Transitions are as follows:

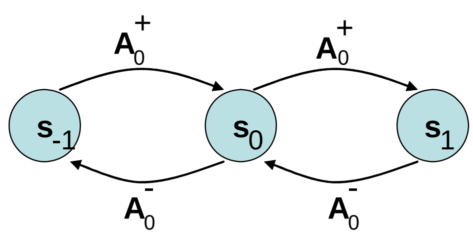

Transition function

As the above figure illustrates, if the agent finds himself in state s_1 when he is called upon to make his decision, there may be many histories that have led to him being in state s_1. We assume that the agent is limited in that he is unable to distinguish more finely between histories. Consequently, S and T are simply devices that generate a particular pooling of the histories that the agent faces when making a decision.

Evidence

Optimal behavior.

Our aim is to understand the consequences of limits on mental states and exogenous transitions among those states. To focus on those aspects of mental processing, we assume that the agent behaves optimally, contingent on the state he is in. We do not claim that there are not additional biases in the way agents process the information they receive; indeed, there is substantial work investigating whether, and how, agents may systematically manipulate the information available to them.[5]

Formally, the set of mental states S, the initial state, the transition function T and Nature's choice of the true state generate a probability distribution over the mental states the agent will be in in any period t that he might be called

[5]Papers in this area in psychology include Festinger (1957), Josephs et *al.* (1992) and Sedikides et *al.* (2004). Papers in economics that pursue this theme include Benabou and Tirole (2002, 2004), Brunnermeier and Parker (2005), Compte and Postlewaite (2004), and Hvide (2002).

upon to make a decision. These distributions along with the probability distribution over the periods that he must make a decision, determine a probability distribution over the state the agent will be in when he makes a decision, $q^{S,T}(s)$. This along with the probability distribution over the true state θ determines a joint distribution over (s, θ), denoted $q^{S,T}(s, \theta)$.

We assume that (possibly through experience) the agent is able to identify the optimal decision rule $a(s, c)$[6]; that is, that the agent can maximize:

$$\sum_{s,\theta} q^{S,T}(s, \theta) E_c u(a(s, c), \theta, c).$$

Call $\pi(s) = \Pr\{\theta = 1 \mid s\}$, the Bayesian updated belief. The expected utility above is maximal when the agent chooses $a = 1$, when $\pi(s) - c_1 \geq 1 - \pi(s) - c_2$, and the maximum expected utility can be rewritten as

$$V(S, T) = \sum_{s} q^{S,T}(s) v(\pi(s)).$$

We illustrate our approach next with specific examples.

Computations.

We will illustrate how one computes the distribution over states prior to decision making, conditional on the true state being $\theta = 1$. Define

$$p = \Pr\{\widetilde{l} > 1 \mid \theta = 1\}.$$

Suppose the true state is $\theta = 1$. Denote by ϕ the distribution over states of mind at the time the agent makes a decision[7]:

$$\phi = \begin{pmatrix} \phi(s_1) \\ \phi(s_0) \\ \phi(s_{-1}) \end{pmatrix}.$$

Also denote by ϕ^0 the initial distribution over states of mind (i.e., that puts all weight on s_0, so that $\phi^0(s_0) = 1$); Conditional on $\theta = 1$, one additional signal moves the distribution over states of mind from ϕ to $M\phi$, where

$$M = \begin{pmatrix} p & p & 0 \\ 1-p & 0 & p \\ 0 & 1-p & 1-p \end{pmatrix},$$

is the transition matrix associated with the mental process (S, T).

[6] It is indeed a strong assumption that the agent can identify the optimal decision rule. As stated above, our aim is to demonstrate that even with the heroic assumption that the agent can do this, there will he will systematically make mistakes in some problems.

[7] ϕ is the conditional distribution on S given $\theta = 1$, that is, $\phi(s) = q^{S,T}(s|\theta = 1)$.

Starting from ϕ^0, if the true state is $\theta = 1$, then the distribution over states of mind at the time the agent takes a decision will be:

$$\phi = (1-\lambda)\sum_{n\geq 0}\lambda^n (M)^{n+1}\phi_0$$

or equivalently,

$$\phi = (1-\lambda)(I-\lambda M)^{-1}M\phi^0. \tag{1}$$

These expressions can then be used to compute $q^{S,T}(s,\theta)$, hence $\pi(s)$ and $q^{S,T}(s)$. For example, $q^{S,T}(s_1, 1) = \pi_0 \cdot \phi(s_1)$.

More generally, given any mental process (S,T), one can associate a transition matrix M that summarizes how an additional signal changes the distribution over states of mind, and then use (1) to derive $q^{S,T}(s)$, and subsequently, expected welfare $v(S,T)$.

A fully symmetric case.

We return to example 1, and assume that $\pi_0 = 1/2$ and a symmetric signal structure.

Assumption 1: $x \in [0,1]$, $f(x \mid \theta = 1) = f(1-x \mid \theta = 2)$.

Under assumption 1, $\Pr\{l > 1 \mid \theta = 1\} = \Pr\{l < 1 \mid \theta = 2\}$. We let $p = \Pr\{l > 1 \mid \theta = 1\}$. Note that we must have $p > 1/2$. Given these symmetry assumptions, $\pi(s_0) = 1/2$ and $\pi(s_1) = 1 - \pi(s_{-1})$.

The following proposition gives $\pi(s_1)$ and $q^{S,T}(s_1)$ as a function of λ and p.

Proposition 1: $\pi(s_1) = p\frac{1-\lambda(1-p)}{1-2\lambda p(1-p)}$ and $q^{S,T}(s_1) = \frac{1-2\lambda(1-p)p}{2-2\lambda^2(1-p)p}$.

As p or λ get closer to 1, beliefs become more accurate (conditional on s_1 or s_{-1}), and there is a greater chance that the agent will end up away from s_0. At the limit where p is close to 1, the agent almost perfectly learns the correct state before taking a decision. Note that $\pi(s_1) > p$ for all values of λ. Intuitively, being at state of mind $s = s_1$ means that the balance of news in favor/against state s_1 tilts in favor of s_1 by on average more than just one signal. The reason is that if the agent is in state 1, it is because he just received a good signal, and because last period he was either in state 0 (in which case, by symmetry, the balance must be 0) or in state 1 (in which case the balance was already favorable to state s_1).

Finally, note that proposition 1 permits welfare analysis. To get a simple expression, assume that the distribution over costs (c_1, c_2) is symmetric, which implies that $v(\pi) = v(1-\pi)$. Then expected welfare is:

$$2q^{S,T}(s_1)v(\pi(s_1)) + (1 - 2q^{S,T}(s_1))v(\frac{1}{2}).$$

As one expects, welfare increases with the precision of the signal (p).

3 Comparing mental processes

Our objective. Our view is that a mental processing system should work well in a variety of situations, and our main interest lies in understanding which mental process (S, T) works reasonably well, or better than others. In this section, we show that there is always a welfare gain to ignoring mildly informative signals.

An improved mental process: ignore mild evidence

We return to example 1, but we now assume that a signal must be minimally informative to generate a transition, that is, to be taken as evidence for or against a particular state. Formally, we define:

$$A^+ = \{\widetilde{l} > 1 + \beta\} \text{ and } A^- = \{\widetilde{l} < \frac{1}{1+\beta}\}.$$

In other words, the event $A = \{\frac{1}{1+\beta} < \widetilde{l} < 1 + \beta\}$ does not generate any transition. Call (S, T^β) the mental process associated with these transitions. Compared to the previous case ($\beta = 0$), the pooling of histories is modified. We may expect that because only more informative events are considered, beliefs conditional on s_1 or s_{-1} are more accurate. However, since the agent is less likely to experience transitions from one state to another, the agent may have a greater chance of being in state s_0 when making a decision.

The symmetric case

We return to the symmetric case, and assume that costs are drawn uniformly. We wish to plot expected welfare as a function of β for various values of λ. To this end, we further specify a distribution over signals: we assume that $x \in [0, 1]$ and $f(x \mid \theta = 1) = 2x$.[8]

Each mental process (S, T_β) and state θ generates transitions over states as a function of the signal x. Specifically, let $\alpha = \frac{\beta}{2(1+\beta)}$. When for example the current state is s_0 and x is received, the agent moves to state s_1 if $x > \frac{1}{2} + \alpha$, he moves to state s_2 if $x < \frac{1}{2} - \alpha$, and he remains in s_0 otherwise. Denote by $y = \Pr\{\widetilde{l} \in (1, 1 + \beta) \mid \theta = 1\}$ and $z = \Pr\{\widetilde{l} \in (\frac{1}{1+\beta}, 1) \mid \theta = 1\}$.[9] Conditional on state $\theta = 1$, the transition matrix is given by:

$$M_\beta^{\theta=1} = \begin{pmatrix} p+z & p-y & 0 \\ 1-p-z & y+z & p-y \\ 0 & 1-p-z & 1-p+y \end{pmatrix}$$

As before, this matrix can be used to compute beliefs $\pi(s_1)$ and the probability to reach s_1, that is $q(s_1)$. Increasing β typically raises $\pi(s_1)$ (which is good for welfare), but it makes it more likely to end up in s_0 (which adversely affects welfare). The following graph shows how welfare varies as a function of β for

[8] Given our symmetry assumption, this implies $f(x \mid \theta = 2) = 2(1 - x)$. So signals x above $1/2$ are evidence in favor of $\theta = 1$, and the strength of the evidence ($l(x) = \frac{x}{1-x}$) gets large when x gets close to 1.

[9] $y = (1/2 + \alpha)^2 - (1/2)^2 = \alpha(1 + \alpha)$, and $z = (1/2)^2 - (1/2 - \alpha)^2 = \alpha(1 - \alpha)$.

two values of λ, $\lambda = 0.5$ and $\lambda = 0.8$.

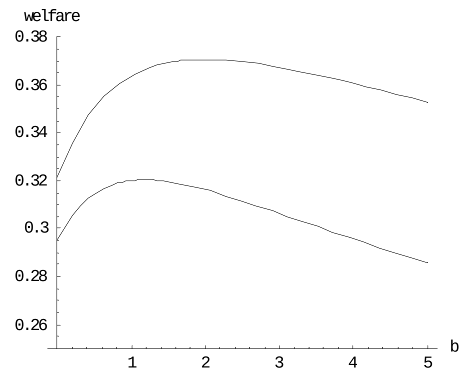

Welfare as a function of β

Note that for a fixed value of λ, for very high values of β there would be little chance of ever transiting to either s_1 or s_2, hence, with high probability the decision would be taken in state s_0. This clearly cannot be optimal, so β cannot be too large. The graph also suggests that a value of β set too close to 0 would not optimal either. The graph illustrates the basic trade-off: large $\beta's$ run the risk of never leaving the initial state while small $\beta's$ have the agent leaving the "correct" state too easily. When λ is sufficiently large, the first effect is small; consequently, the larger λ, the larger is the optimal value of β.

The general case

The advantage of ignoring weakly informative signals in the example above does not depend on the symmetry assumption, nor on the simple mental system of the example. The next proposition states that for any number of states and for any transition function T, and for any value of λ, an agent strictly benefits from having a mental process that ignores poorly informative signals.

Formally, denote by $W(\lambda, \beta)$ the welfare associated with mental process (S, T^{β}). We have:

Proposition 2: There exists $a > 0$ and $\beta_0 > 0$ such that, for all λ and $\beta \in [0, \beta_0]$:

$$W(\lambda, \beta) \geq W(\lambda, 0) + a\beta.$$

Proof: See appendix.

3.1 Consequences of ignoring weak evidence

What are the consequences of the fact that agents would have $\beta > 0$. As we discussed in the introduction, our view is that there is a mental process that summarizes the signals an agent has received in a mental state, and that the agent chooses the optimal action given his mental state when he is called upon to make a decision. Agents do not have a different mental process for every possible decision they might some day face. Rather, the mental process that aggregates and summarizes their information is employed for a variety of problems with different signal structures f. β is set optimally *across* problems, not just for a specific distribution over signals f (that in addition would be correctly perceived). When β is set optimally across problems, it means that for some problems, where players frequently receive mild evidence and occasionally strong evidence, there is a bias towards the theory that generates occasional strong evidence, when Bayesian updating might have supported the alternative theory.

4 Handling perception errors

4.1 Robustness.

Our main result is robust to perception errors. Intuitively, our key result stems from the following observation. Consider histories of signals leading to state s under (S, T^β). Among those histories, some of them would lead to a different state, say s', under (S, T). However, for β small, the only reason why the two final states differ is that one perception $\widetilde{l}$ fell in the interval $(\frac{1}{1+\beta}, 1+\beta)$ and that it thus triggered a different transition under the two mental processes. The problem is that there is a gap between the belief the agent holds at s and s', and thus, a perception $\widetilde{l}$ that is a poorly informative perception optimally would not trigger a change in belief unless the change in belief is very small.

When perceptions are correct, a perception $\widetilde{l}$ that falls in the interval $(\frac{1}{1+\beta}, 1+\beta)$ indeed corresponds to a very poorly informative signal.

When perceptions are not correct in the sense defined earlier, $\widetilde{l}$falls in the interval $(\frac{1}{1+\beta}, 1+\beta)$ either because the true l is close to one and μ is not very large, or because μ is very small. To the extent that μ is drawn from a distribution with a continuous density, for β small enough, the event $\widetilde{l} \in (\frac{1}{1+\beta}, 1+\beta)$ must also be poorly informative about the true state, and our main result extends.

4.2 Comparison with the Bayesian approach

In our approach, the agent receives signals (x) that he tries to interpret, possibly in light of some prior knowledge about the chance of receiving x conditional on θ. This cognitive process generates a perception $\widetilde{l}$ that triggers (or not) a transition in his state of mind. We formalize these processes by defining a distribution over perceptions conditional on each signal received, and a transition function. However we do not have in mind that the agent knows the process that generates perception. The transition function or mental system aggregates these perceptions over time, and the agents choice is only to choose the optimal behavior conditional for each state he may be in. The transition function is *not* optimized for a specific problem, but rather is assumed to be optimal for a set of problems that typically differ in the probability distributions over the signals conditional on the true state of nature and in the expected number of signals the agent will receive prior to making a decision. Even with few states, this is already a difficult task. Our interest then is to investigate modifications of the mental system that increase the agents' welfare across over the set of problems he faces.

In comparison, a fully Bayesian agent would have to perform one of the following tasks. He could know the distribution over signals and errors (where could this knowledge come from?) and invert it to compute posteriors correctly. Alternatively, he would have to find (through experience?) optimal behavior as a function of all possible histories of signals, and for each problem that he faces.

Of course, by adding many states and by looking at the transitions (which map states and perceptions to states) which maximize expected welfare for a given problem, one would expect to eventually approximate the decisions that a fully Bayesian agent would make. But given that the number of possible transitions grows exponentially with the number of states, this illustrates how difficult it is to approximate such behavior.

However complex it is to approximate Bayesian behavior, the problem is an order of magnitude more difficult if the transitions must be derived independently for each problem (characterized in our setting by a distribution over signals f and a value of λ) that one faces. Our view is that while deriving optimal behavior at each state could be problem specific, modifications of transitions (is. changes in β) should be considered to the extent that they generate welfare increases on average across problems. In other words, we have in mind a hierarchical model of behavior in which learning about optimal behavior is not done at the same pace as learning about how to aggregate information.

4.3 Comparing mental processes (continued)

In Section 3, we analyzed mental processes that ignore weak evidence. We now introduce another class of mental processes that attempt to distinguish between strong and mild evidence, and that allow the agent to move directly from one extreme state to the other.

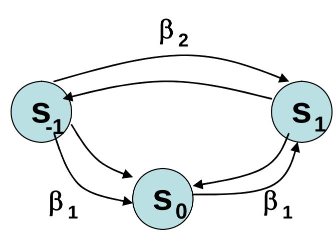

Our motivation is to assess the benefits of allowing for such transitions, and to point out reasons as to why evolutionary forces towards such more complex mental system may be weak. There are indeed at least two difficulties.

(1) we want the mental process to be NOT problem specific. So what is an appropriate threshold for strong evidence is an issue, as it must be appropriate

across problems. For some problems, almost everything might then seem to be strong evidence.

(2) If there are perception errors in strength of evidence, then a more sophisticated mental process might work less well because it may become in effect a two state system. Also even small asymmetries in the distribution over likelihood ratios may then generate biases.

5 The case of many theories

To be added.

6 Related literature

The central theme of this paper – that a decision maker uses a single decision protocol – for a number of similar but not identical problems. The history of this idea in economics goes back at least to Baumol and Quandt (1964). A more systematic modelling of the idea can be found in Rosenthal (1993), where it is assumed that an agent will chose among costly rules of thumb that he will employ in the set of games they face. Lipman (1995) provides a very nice review of the literature on modelling bounds on rationality resulting from limits on agents' ability to process information.

TO BE COMPLETED

7 Appendix

Proof of Proposition 2:

Proof: In what follows, the initial state is assumed to be s_0. However, the analysis below holds for any given initial state, possibly different from s_0. A history h refers to a termination date $\tau \geq 1$ and a sequence of perceived likelihood ratios $(\widetilde{l}_0, \widetilde{l}_1, ..., \widetilde{l}_{\tau-1})$. Denote by H_s^β the set of histories h that lead to state s (starting from initial state s_0) when the mental process is (S, T^β). Also denote by $H_{s,s'}$ the set of histories that lead to state s under (S, T^0) and to state s' under (S, T^β).

For any set of histories H, we denote by $\pi(H)$ the Bayesian posterior, conditional on the event $\{h \in H\}$:

$$\pi(H) = \frac{\Pr(H \mid \theta = 1)}{\Pr(H)} \pi_0$$

By definition of $W(\lambda, \beta)$, we have:

$$W(\lambda, \beta) = \sum_{s \in S} \Pr(H_s^\beta) v(\pi(H_s^\beta))$$

Now let $\bar{W}$ denote the welfare that the agent would obtain if he could distinguish between the sets $H_{s,s'}$ for all s, s'. We have:

$$\bar{W} = \sum_{s,s' \in S} \Pr(H_{s,s'}) v(\pi(H_{s,s'}))$$

We will show that there exist constants c and c' such that

$$W(\lambda, \beta) \geq \bar{W} - c\beta^2 \tag{2}$$

and

$$W(\lambda, 0) \leq \bar{W} - c'\beta. \tag{3}$$

Intuitively, welfare is $\bar{W}$ when the agent can distinguish between all $H_{s,s'}$. Under (S, T^β), he cannot distinguish between all $H_{s,s'}$:

$$H_s^\beta = \cup_{s'} H_{s',s}$$

Under (S, T^0), he cannot distinguish between all $H_{s,s'}$ either, but the partitioning is different:

$$H_s^0 = \cup_{s'} H_{s,s'} \tag{4}$$

Because each mental process corresponds to a pooling of histories courser than $H_{s,s'}$, and because v is a convex function, both $W(\lambda, 0)$ and $W(\lambda, \beta)$ are smaller than $\bar{W}$. What we show below however is that the loss is negligible in latter case (of second order in β), while it is of first order in the former case.

We use three Lemmas, the proofs of which are straightforward.

Lemma 1: There exists a constant c such that:

$$\left|\pi(H_{s',s}) - \pi(H_s^\beta)\right| \leq c\beta. \tag{5}$$

Lemma 2: For any s, there exists $s' \neq s$ with $\Pr(H_{s,s'}) = O(\beta)$ such that :

$$\left|\pi(H_{s,s'}) - \pi(H_s^0)\right| \geq c'.$$

Lemma 3: Let $m, \bar{m}$ such that $\bar{m} \geq v'' \geq m$. For α small, we have, forgetting second order terms in α:

$$\alpha \frac{\bar{m}}{2}(\pi^1 - \pi^0)^2 \geq (1-\alpha)v(\pi^0) + \alpha v(\pi^1) - v(\alpha\pi^0 + (1-\alpha)\pi^1) \geq \frac{m}{2}(\pi^1 - \pi^0)^2.$$

Since

$$\sum_{s \in S} \Pr(H_{s,s'})\pi(H_{s,s'}) = \pi(\bigcup_s H_{s,s'}) \sum_{s \in S} \Pr(H_{s,s'}) = \pi(H_{s'}^\beta) \Pr(H_{s'}^\beta),$$

we have:

$$\begin{aligned}\bar{W} - W(\lambda,\beta) &= \sum_{s'\in S}\left[\sum_{s\in S}\Pr(H_{s,s'})[v(\pi(H_{s,s'})) - v(\pi(H^{\beta}_{s'}))]\right] \\ &= \sum_{s'\in S}\left[\sum_{s\in S}\Pr(H_{s,s'})[v(\pi(H_{s,s'})) - v(\sum_{s\in S}\Pr(H_{s,s'})\pi(H_{s,s'}))]\right].\end{aligned}$$

Applying Lemma 3 thus yields:

$$\bar{W} - W(\lambda,\beta) \leq c \max_{s,\widehat{s},s'} \mid \pi(H_{s,s'}) - \pi(H_{\widehat{s},s'}) \mid^2,$$

and Lemma 1 gives a lower bound on $W(\lambda,\beta)$.

To get the upper bound on $W(\lambda,0)$, we use:

$$\sum_{s'\in S}\Pr(H_{s,s'})\pi(H_{s,s'}) = \pi(\bigcup_{s'} H_{s,s'}) \sum_{s'\in S}\Pr(H_{s,s'}) = \pi(H^0_s)\Pr(H^0_s),$$

and write:

$$\bar{W} - W(\lambda,\beta) = \sum_{s\in S}\left[\sum_{s'\in S}\Pr(H_{s,s'})[v(\pi(H_{s,s'})) - v(\pi(H^0_s))]\right].$$

Lemmas 2 and 3 (using s and s' as defined in Lemma 2) then yield a lower bound:

$$\bar{W} - W(\lambda,\beta) \geq c\beta[\pi(H_{s,s'}) - \pi(H^0_s)]^2.$$

8 Bibliography

Baumol, William J. and Richard E. Quandt (1964), "Rules of Thumb and Optimally Imperfect Decisions," *American Economic Review* 54(2), pp 23-46.

Benabou, R. and J. Tirole (2002). "Self-Confidence and Personal Motivation," *Quarterly Journal of Economics*, 117(3), 871-915.

Benabou, R. and J. Tirole (2004). "Willpower and Personal Rules," *Journal of Political Economy*, 112 (4), 848-886.

Charness, Gary and Dan Levin (2005), "When Optimal Choices Feel Wrong: A Laboratory Study of Bayesian Updating, Complexity, and Affect," *American Economic Review*, 96, 1300-1309.

Compte, O., and Andrew Postlewaite. (2004) "Confidence-Enhanced Performance," *American Economic Review*, 94, 1536-1557.

Compte, O. and A. Postlewaite (2009a), "Effecting Cooperation", mimeo, University of Pennsylvania.

Compte, O. and A. Postlewaite (2009b), "Repeated Relationships with Limits on Information Processing", mimeo, University of Pennsylvania.

Cover, T. and M. Hellman (1970), "Learning with Finite Memory," *Annals of Mathematical Statistics* 41, 765-782.

Dow, J. (1991) "Search Decisions with Limited Memory," *Review of Economic Studies* 58, 1-14.

Festinger, L. (1957), *A theory of cognitive dissonance*. Stanford, CA: Stanford University Press.

Gilovich, T. (1991), *How We Know What Isn't So.* New York, NY: The Free Press.

Griffin, D. and A. Tversky (1992) "The Weighing of Evidence and the Determinants of Confidence," *Cognitive Psychology* 24, 411-435.

Hellman, M. E. and T. M. Cover (1973). "A Review of Recent Results on Learning with Finite Memory," *Problems of Control and Information Theory*, 221-227.

Hvide, H. (2002). "Pragmatic beliefs and overconfidence," *Journal of Economic Behavior and Organization*, 2002, 48, 15-28.

Jenkins, H.M. and W.C. Ward (19650 "Judgement of contingency between responses and outcomes," *Psychological Monographs*

Josephs, R. A., Larrick, R. P., Steele, C. M., and Nisbett, R. E. (1992). "Protecting the Self from the Negative Consequences of Risky Decisions," *Journal of Personality and Social Psychology*, 62, 26-37.

Kahneman, D., P. Slovic, and A. Tversky (1982) *Judgement under Uncertainty: Heuristics and Biases,* New York, NY: Cambridge University Press.

Lamb, Emmet J. and Sue Leurgans (1979), "Does Adoption Affect Subsequent Fertility?", *American Journal of Obstetrics and Gynecology* 134, 138-144.

Lichtenstein, S. and Fischhoff, B. (1977) "Do those who know more also know more about how much they know? The calibration of probability judgements," *Organizational Behavior and Human Performance*, 20, 159-183.

Lipman, Barton (1995), "Information Processing and Bounded Rationality: A Survey," *The Canadian Journal of Economics* 28(1), pp. 42-67.

Rosenthal, Robert W., 1993. "Rules of thumb in games," Journal of Economic Behavior & Organization, Elsevier, vol. 22(1), pages 1-13, September.

Sedikides, C., Green, J. D., and Pinter, B. T. (2004). "Self-protective memory," In: D. Beike, J. Lampinen, and D. Behrend, eds., *The self and memory.* Philadelphia: Psychology Press.

Ward, W. C. and H.M. Jenkins (1965). "The display of information and the judgment of contingency," *Canadian Journal of Psychology*

Wilson, A. (2004), "Bounded Memory and Biases in Information Processing," mimeo, University of Chicago.