

Self-esteem preferences and imaginary effects in information processing

Wojciech Olszewski*

September 18, 2016

1 Introduction

Recently, a number of behavioral phenomena, such as polarization or overconfidence, have been widely documented. For example, a classic example of experimental evidence for polarization is Darley and Gross (1983) who randomize subjects into different groups. They show one group evidence suggesting a child is from a high socioeconomic background; another that she is from a low socioeconomic background. The former predict the child's reading abilities are higher than the latter. The groups then watch a film of the child taking an oral test on which she answers some questions correctly and others incorrectly. Those who received the information that the child came from a high socioeconomic background, rate her abilities higher than before; those who received the information indicating she came from a low socioeconomic background rate her abilities lower than before.

These phenomena are inconsistent with the traditional Bayesian modelling of information and its processing. The traditional modelling assumes that people act as if their information could be summarized by probabilistic beliefs, which are updated by the Bayes rule when new information arrives. This implies that both groups in the Darley and Gross experiment should update their belief for more positive or more negative, or their assessments should become closer.

Numerous explanations of these phenomena have been offered, ranging from questioning the inconsistency of some of these phenomena with the traditional models to proposing a collection of new behavioral models obtained by replacing the rationality assumption with some bounded-rationality assumption that resonates well with the author's intuition.¹ These and other behavioral phenomena, and their implications on economic modelling (including information processing) were the object of perhaps the hottest discussion in economic theory in recent years. Perhaps, the most typical opinion would be that behavioral phenomena

*The author is grateful to Piotr Dworzak and Ran Spiegler for helpful comments and suggestions.

¹We discuss these explanations in the section on the existing related literature.

are inconsistent with traditional modelling, but the recently proposed behavioral approach suffers from not imposing any discipline on the modeler.

In this paper, we suggest a theory which generalizes the traditional approach, departs from the traditional approach in a very restricted way, and captures a quite large number of the behavioral phenomena concerning information processing. The general premise of our theory is based on two components: (1) in practice, people often receive vague signals, which leaves room for interpretation; (2) they interpret vague signals in the way which is convenient to them. More specifically, the interpretation not only affects their direct payoffs, but often also reveals what they think about themselves, or affect the way they wish to view themselves (or be viewed by others). This approach is related to self-enhancement theory in psychology literature (see Greenwald, 1980). According to the self-enhancement theory people have a desire “to see themselves favorably as competent human beings,” which increases their feelings of personal satisfaction and worth. Therefore, they distort information processing so as to select, interpret and recall information in a way that supports a positive self-image.

More specifically, we assume in this paper that agents enjoy thinking of their ability (of processing information) as being high. They receive information in the form of, possibly vague, exogenous signals, modelled as sets of probability distributions over signal values, contingent on the state of the world. Agents can (or even must) interpret any vagueness of information, that is, choose their beliefs, which enables them to enhance their self-esteem. However, this choice of beliefs and self-esteem is limited in two ways. First, the signals may not be too vague, that is, may consist of similar probability distributions. For example, when the information has the form of single probability distributions, agents’ beliefs are uniquely determined by the information they receive, exactly as in the traditional Bayesian analysis. We emphasize that the signals that agents receive are exogenous in our model; agents cannot choose the level vagueness of their information. We will discuss the choice of signals in the final section of the paper.

In addition, agents interpret any vagueness only when they must take actions. One may say that these are agents’ actions that reveal their true beliefs. We emphasize the interpretation that agents’ beliefs about their ability change only with their actions, and only their actions reveal or admit what they really think about their ability. It cannot happen that agents’ beliefs about ability are inconsistent with their actions, given the information they have. This imposes some endogenous limits on the agents’ self-esteem.

So, agents face a trade-off between direct payoff implications of their actions, and what they reveal or admit about their ability. We demonstrate that our theory generates what we call *imaginary confirming effect*. This is a processing of information biased towards prior views, caused by a preference for high assessment of own ability. This effect is obviously a potential source of polarization. We also demonstrate that our theory generates another, a sort of perverse effect, which we call *imaginary discrediting effect*. This is a processing of information biased against prior views. This effect is not related to as widely documented phenomena as polarization, but seems to be commonly present in everyday life. For example, academics often interpret various evaluation letters referring to what the authors of these letters were trying to say

between the lines. This kind of excessive diligent reading is precisely what the new effect is trying to capture.

We identify situations in which the imaginary effects in information processing are present, and are likely to be stronger. Although most of the analysis is conducted in a two-period version of our model, we also explore whether the agents in our model will eventually learn about the true state of the world, when they receive in every period a signal which contains a “grain of truth.” We provide an example showing that this may not happen. In particular, agents’ opinions may diverge over time. We show that agents learn about the state of the world, and their opinions must converge over time, only if sufficiently many of their signals contain a “sufficiently large grain of truth.” We conclude the paper with some additional examples of insights offered by our model regarding behavioral phenomena in information processing.

Related Literature

Several motives which our theory refers to have been explored in some earlier papers. The use of sets, as opposed to single objects, as a representation of vagueness of information dates from Arrow and Hurwicz (1972). In their model, a decision maker chooses a set of objects, and is completely ignorant regarding which object from the chosen set she will actually “consume.” More recent papers by Ahn (2008), Olszewski (2007) and Vierø (2009) suggest to use sets of probability distributions in the place of single probability distributions to model ambiguity or vagueness of information. In the present paper, we apply the approach suggested in these latter papers. Our paper is not the first which includes self-assessment into utility function. For example, Zábajnik (2004) explores a model in which individuals can learn about their abilities, and where an individual’s utility is a function of self-assessments. Individuals in this model keep testing their abilities until their self-assessments become favorable enough, at which point they stop. Consequently, an excessively large share of the population ends up with a high opinion about their abilities. The maintenance and enhancement of self-esteem is also emphasized in Bénabou and Tirole (2002), who study a model in which individuals can, within some limits, affect the probability of remembering a given piece of data. In deciding whether to try to repress bad news, the individual weighs the benefits from preserving his effort motivation (enhanced by high self-esteem) against the risk of becoming wrongly optimistic.

Finally, there exist models in which agents have some freedom in choosing their beliefs. Among them, Brunnermeier and Parker (2005) seems to be closest to our paper. In their model, agents first choose their beliefs, and then taking their beliefs as given, they choose control variables and the implied evolution of state variables to maximize their utility. Agents may, for example, overestimate the probability of good outcomes, which will make them happier, but distorted beliefs distort actions, and worsen physical outcomes. A similar trade-off is faced by agents in our model, where they overestimate their ability, which makes them happier, but this distorts their actions and worsens physical outcomes. An important advantage of our model is that it allows for studying information acquisition (see Section 7 for details.)

One of the applications of our model concerns polarization. There is now a substantial literature attempting to explain the phenomenon of polarization. Andreoni and Mylovannov (2012) and Kondor (2012) show that polarization can be generated within standard economic models under certain signal structures.

Acemoglu, Chernozhukov and Yildiz (2009) show that if agents have different priors on some payoff-relevant parameters, and also different distributions on signals conditional on the parameters, then the posteriors on parameters can diverge, even after receiving infinite sequences of signals. Zimmer and Ludwig (2009) and Baliga, Hanany and Klibanoff (2013) relate polarization to ambiguity. In the former case, polarization occurs because agents happen to apply different updating rules in the presence of non-additive beliefs. In the latter case, polarization is implied by ambiguity aversion. Dunning (1989), Van den Steen (2004), and Santos-Pinto and Sobel (2005) argue that people may have varying notions of, say, reading ability. The last two papers argue further that as a result of different notions, people invest in skills in different ways.

Our explanation of polarization is probably closest to the Rabin and Schrag (1999) model of confirmatory bias in which agents may misread a signal conflicting with what they believe, and thus updating is simply assumed to be biased in the direction of current beliefs, directly generating polarization. An important new feature of our model is that it provides discipline for agents' subjective beliefs. Agents may choose their beliefs only when their information is vague, and this choice is limited exogenously as the information they receive is typically not entirely vague. This choice is also limited endogenously as biased choices distort actions and generate inferior outcomes.

Our analysis is clearly related to the concept of overconfidence, since the self-esteem effects amplify agents' belief in their own ability. A classic example of experimental evidence on overconfidence is so called better-than-median effect. This effect has been documented, for example, in a study Svenson (1981) who found that 77% of Swedish subjects felt they were safer drivers than the median, and 69% felt they were more skillful. There is now a large literature on overconfidence, and summarizing it would not be particularly useful for the present paper. A detailed overview of the literature can, for example, be found in Benoit and Dubra (2011), who argue that lots of the evidence on overconfidence can be explained within the traditional rational setting. In turn, most of this literature argues that agents' subjective beliefs must differ from the belief that agents would have if they were processing information rationally.

Our model can be placed between the rational approach (like that in Benoit and Dubra), and more common approach assuming incorrect beliefs. Agents' beliefs in our model are distorted by self-esteem effects, but they cannot be described as incorrect, since they are consistent with the vague information that the agents have. We are not explaining any evidence on overconfidence in this paper, since we prefer to perform the analysis in a single-agent setting, and not explore any the issues concerning comparisons across individuals. We wish to emphasize, however, that our setting provides (in Section 5) an original way of viewing the empirical evidence that the better-than-median effect becomes smaller over time, or vanishes with accumulated information or experience. For example, Benoit and Dubra cite various evidence from interviewing professional truck drivers, and according to this evidence, younger drivers rate themselves as being better overall drivers than their peers, while older drivers rank themselves comparably to drivers in their age group.

2 Examples

Example 1 (polarization). In this example, we provide our explanation of the phenomenon of polarization. Consider two agents, named T and F. (The agents will not interact, so the analysis will be conducted within a decision-theoretic framework.) Each agent has the fifty-fifty independent priors that some claim is true or false, and that she is smart or dumb. In each of two periods, the agents observe the realization (value) of some signal that is relevant for determining the correctness of the claim, as well as for learning about the agent’s ability. Each signal takes one of two possible values: ‘true’ or ‘false.’ If the agent is smart, the value of a signal is equal to the logical value of the discussed claim with probability $3/4$; if the agent is dumb, the value is equal to the logical value of the discussed claim only with probability $1/2$.

Suppose that in period 1 each agent privately observed² a realization of signal with this structure. The signals of the two agents are conditionally independent. Suppose that the value observed by agent T was ‘true,’ while the value observed by agent F was ‘false.’ Suppose finally that in period 2 the realization of a public vague signal is observed. The public signal has the same structure as the signals in period 1, but suppose that the agents cannot tell what was the signal’s value, whether it was ‘true’ or it was ‘false,’ even they cannot tell what is the chance of the signal taking each of the two values.³ Note that the signal in period 2 is the same for both logical values of the claim, so the signal contains no information about the state of the world. (This is obviously an extremely vague signal.)

Consider first the beliefs of the two agents after observing the signal in period 1. Updating her prior by the Bayes rule, agent T believes that the claim is true with probability

$$\text{prob}(\text{smart}) \cdot \text{prob}(\text{‘true’} \mid \text{smart, true}) + \text{prob}(\text{dumb}) \cdot \text{prob}(\text{‘true’} \mid \text{dumb, true}) = \frac{1}{2} \cdot \frac{3}{4} + \frac{1}{2} \cdot \frac{1}{2} = \frac{5}{8},$$

and that she is smart with probability $1/2$. Similarly, agent F assigns a probability of $5/8$ to the claim being false, and a probability of $1/2$ to being smart.

Now, consider the beliefs of agent T after obtaining the vague signal in period 2. If agent T assumes that its value was ‘true,’ the chance observing two values ‘true’ contingent on each of the four possible states of the world is:

the state of the world	the chance of values ‘true’, ‘true’
true, smart	$\frac{3}{4} \cdot \frac{3}{4} = \frac{9}{16}$
true, dumb	$\frac{1}{2} \cdot \frac{1}{2} = \frac{4}{16}$
false, smart	$\frac{1}{4} \cdot \frac{1}{4} = \frac{1}{16}$
false, dumb	$\frac{1}{2} \cdot \frac{1}{2} = \frac{4}{16}$

²Since the analysis is conducted within a decision-theoretic framework, it does not matter whether a signal is private or public. It will only matter when we interpret the example as an explanation of polarization.

³In our model, agents will actually see the values of signals, but vague signals will have the form of sets of lotteries. The signal from this example can be represented as a set of two lotteries: one having the structure as described, and one taking the logical value of the claim with probability $1/4$ when the agent is smart, and taking the logical value of the claim with probability $1/2$ when the agent is dumb.

Thus, the agent assigns a probability of $13/18$ to the claim being true, and a probability of $5/9$ to her being smart. If agent T assumes that the value was ‘false,’ the chance of observing one signal ‘true’ and one signal ‘false’ contingent on each of the four possible states of the world is:

the state of the world	the chance of values ‘true’, ‘false’
true, smart	$\frac{3}{4} \cdot \frac{1}{4} = \frac{3}{16}$
true, dumb	$\frac{1}{2} \cdot \frac{1}{2} = \frac{4}{16}$
false, smart	$\frac{1}{4} \cdot \frac{3}{4} = \frac{3}{16}$
false, dumb	$\frac{1}{2} \cdot \frac{1}{2} = \frac{4}{16}$

Thus, the agent assigns a probability of $1/2$ to the claim being true, and a probability of $3/7$ to her being smart.

In our model, agent T will have the two posterior beliefs over the claim and her ability. These two posteriors induce two marginal beliefs over the claim being true, which are $13/18$ and $1/2$. When the agent is supposed to give a single probability assessment of the claim being true, she must make the choice how much weight to assign to each posterior. In our model, the agent is indeed assumed to choose from the convex hull of her beliefs. The agent’s choice of her assessment of the claim is not independent of her choice of the assessment of her ability. The choice of one assessment determines the other assessment, because the two choices must be supported by the same weights assigned to the multiple posteriors. This is what we mean when we say that agents’ actions reveal something about their beliefs regarding their ability, or when we say that by taking actions agents admit something about their beliefs regarding their ability.

More precisely, by making the assessment of the claim being true with probability $p \in [1/2, 13/18]$, the agent assesses her being smart with the probability $q \in [3/7, 1/2]$. The left ends of these intervals correspond to weight 1 assigned to posterior obtained by assuming that the value of the vague signal was ‘false,’ and the right ends correspond to weight 1 assigned to posterior obtained by assuming that the value of the vague signal was ‘true’. Therefore, the two assessments are related by the formula

$$q = \frac{4}{7}p + \frac{1}{7}.$$

Suppose first that agents obtain utility only from making the right assessment of the discussed claim, and whenever they have two beliefs they maximize this utility by making the assessment equal to the average of the two. This implies that agent T after obtaining the vague signal revises her assessment of the claim being true to

$$\frac{\frac{13}{18} + \frac{1}{2}}{2} = \frac{11}{18} < \frac{5}{8}.$$

Similarly, agent F revises her assessment of the claim being true to $7/18 > 3/8$. So, the two assessments converge.

Suppose now that agents obtain utility only from their assessment about being smart, and this utility increases with the assessment. Then, agent T will argue that the claim is true with probability $13/18 > 5/8$.

Similarly, agent F will argue that the claim is true with probability $5/18 < 3/8$. This implies that agents polarize, despite observing the same (vague) signal.

Typically, agents care about both: making the right assessment of the discussed claim and the implicitly revealed (or admitted) assessment about being smart. So, their preferences can be represented by indifference curves on the plane, where the former ‘good’ is represented on p -axis and the latter ‘good’ is represented on q -axis. The (dis)utility from the former good is measured by the departure from the average of the two assessments corresponding to posterior beliefs, and the utility of the latter good increases with the assessment. In this case, agents will polarize as long as the marginal rate of substitution between the two goods is sufficiently high. A typical choice problem is illustrated in Figure 1.

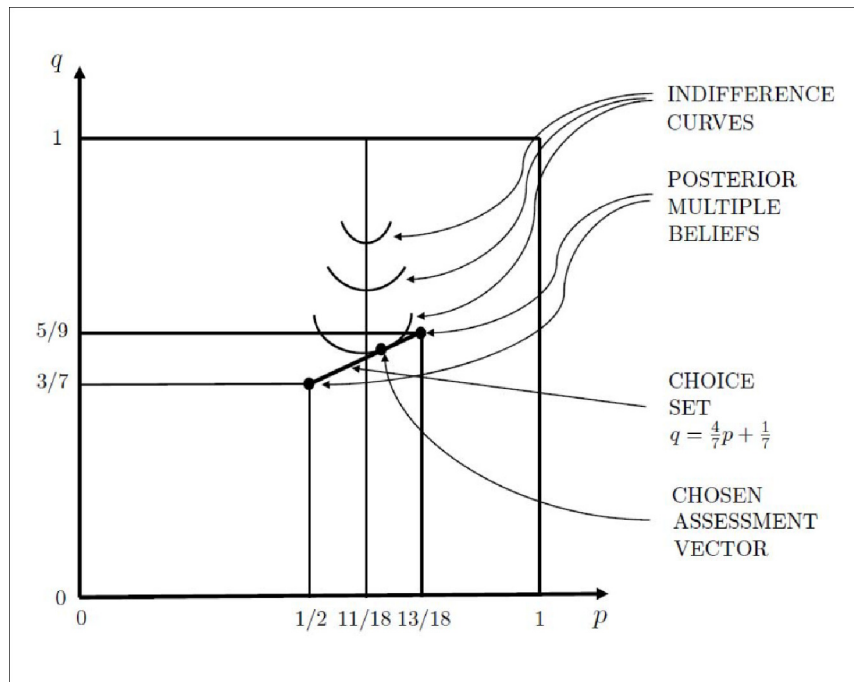


Figure 1. A typical choice problem: indifference curves are horizontal at $p = 11/18$, since this is the agent’s bliss point in her action utility; the indifference curves increase in the direction of q , since the agent’s self-esteem utility increases in q .

Example 2 (‘diligent reading’). Scientists often have to write peer evaluation letters. Despite these letters being confidential, scientists are typically not willing to provide negative opinions about their peers. Negative opinions are often hidden ‘between the lines,’ and positively looking letters are interpreted as bearing negative opinions. In this example, we model this kind of diligent reading, which will imply that there is too much of diligent reading.

As in Example 1, let an agent have fifty-fifty independent priors that some claim is true, and that she is smart. Suppose that in period 1, the agent received a ‘true’ signal about the claim, after which the agent

revised her belief of the claim being true to $5/8$, and did not change the belief about her ability (exactly as in Example 1). Suppose that this signal is conditionally independent of the signal the agent is to receive in period 2. The signal in period 2 is vague. It has two possible interpretations. Under the first interpretation, the signal contains unambiguous information. More specifically, with probability $3/4$ the value of this signal is equal to the logical value of the claim, no matter whether the agent is smart or dumb. Under the other interpretation, the signal contains negative information, but to have a chance of extracting it, the agent must be smart, to be able to "read what is written between lines". More specifically, the signal is fifty-fifty contingent on the agent being dumb, both when the claim is true and when it is false; it is also fifty-fifty when the agent is smart but the claim is true, but the chance of the signal 'true' is $3/4$ (so the chance of the signal 'false' is $1/4$) contingent on the agent being smart when the claim being false. So, reading between the lines is modelled by value 'true' containing a negative message.

Suppose that the value of the signal that the agent observes in period 2 is 'true'. Interpreting this signal as containing unambiguous information, the agent updates in the Bayesian way its prior about the claim and her ability to

$$\frac{\frac{5}{8} \cdot \frac{3}{4}}{\frac{5}{8} \cdot \frac{3}{4} + \frac{3}{8} \cdot \frac{1}{4}} = \frac{5}{6}$$

and $1/2$, respectively. And under the other interpretation, the posteriors are

$$\frac{\frac{5}{8} \cdot \frac{1}{2} \cdot \frac{1}{2} + \frac{5}{8} \cdot \frac{1}{2} \cdot \frac{1}{2}}{\frac{5}{8} \cdot \frac{1}{2} \cdot \frac{1}{2} + \frac{5}{8} \cdot \frac{1}{2} \cdot \frac{1}{2} + \frac{3}{8} \cdot \frac{1}{2} \cdot \frac{1}{2} + \frac{3}{8} \cdot \frac{1}{2} \cdot \frac{3}{4}} = \frac{4}{7}$$

and $19/35$.

When the agent's utility is only a function of the correctness of her assessment of the claim, her estimate for the claim being true is $59/84 > 5/8$ (which is the average of $5/6$ and $4/7$). However, this yields only the estimate of $73/140$ of the agent being smart (this is the average of $1/2$ and $19/35$). Instead, by estimating the claim to be true with probability $5/8$, the agent estimates herself to be smart with probability

$$\frac{63}{308} \cdot \frac{1}{2} + \frac{245}{308} \cdot \frac{19}{35} = \frac{47}{88} > \frac{73}{140}.$$

(To see why, notice that $5/8$ is the weighted average of $5/6$ and $4/7$ with weights $63/308$ and $245/308$.) So, if the agent's utility is also a function of revealed self-esteem, the agent may not change her assessment about the claim being true, estimating its probability for $5/8$, but may have a higher belief of her being smart. The agent may even reduce her assessment about the claim being true, reaching even a higher belief of her being smart. As in Example 1, her choice depends on the marginal rate of substitution between the two goods. One way or another, the utility of feeling smart may only reduce the assessment of the correctness of the claim compared to the assessment of the agent who cares only about the correctness.

3 Model

An agent lives in multiple periods, in a world of unknown state. In this paper, we assume that the state has the form $(\omega, \theta) \in \Omega \times \Theta = \{\text{true}, \text{false}\} \times \{\text{smart}, \text{dumb}\}$. We will also continue referring to the claim story

from our examples. The agent has a prior π_0 over $\Omega \times \Theta$. In every period, the agent receives some, possibly vague, signal about the state of the world. The signal in period t is modelled as the realization (value) $x_t \in X$ of a lottery, chosen by Nature contingent on the state of the world from a set of conditional lotteries $L_t \subset \Delta(X) \times \Omega \times \Theta$. (We will call L_t a signal.) We will assume, in this paper, that $X = \{\text{'true,' 'false'}\}$, and that sets L_t are finite. Actually, we will consider only some specific sets L_t , which will be defined later. Our agent knows the set L_t from which this lottery has been chosen, but does not know which specific lottery has been chosen by Nature, neither has any probability distribution over lotteries that is “used” by Nature.

Sets of lotteries are used to capture the vagueness of signals. In practice, many signals are vague. For example, suppose we observe a child answering some questions regarding the text that she has just read, and we are supposed to evaluate the reading skills of this child. Then, the fact that the child answered some questions correctly, and other questions incorrectly places some constraints on the evaluation, but we are typically unable to provide any precise estimate (e.g., compared to the age group of the child). Similarly, by reading in an evaluation letter about some strengths and weaknesses of a researcher, we obtain only vague information about the quality of his or her future research. Of course, sets of lotteries are only one, not the only tool that can be used for capturing information vagueness.

After receiving signal L_t , the agent takes an action p_t . This action has a direct payoff implication, but also affects, in the way that will be described later, an indirect payoff coming from a (*revealed*) *self-esteem*. In the present paper, we assume that the agent’s actions belongs to interval $[0, 1]$, and this action is interpreted as her assessment of ω being equal to 1. She obtains an action flow payoff which is a function of her action and the actual chance of $\omega = 1$. More specifically, if the agent knew that the state ω is equal to 1 with probability p , her payoff from taking action p_t would be a decreasing function u of $|p_t - p|$. The self-esteem flow payoff, in turn, is an increasing function v of the the probability q_t assigned to $\theta = 1$. Since our setting involves vagueness, we must generalize the payoffs to accommodate situations in which the agent’s information I_t is not a single probability distribution over states, but is a set of distributions. This is done by some payoff operators U and V which extend functions u and v from single probability distributions to sets of distributions. For example, the agent’s action payoff may be a function of the distance between her action p_t and the average of the highest and the lowest p over all lotteries in any set. Similarly, the agent’s self-esteem payoff may be the average of the highest and the lowest q over all lotteries in any set. The agent’s overall flow payoff (that is, the payoff she receives in period t) depends on the two payoffs, aggregated by some function ϕ that is increasing in each argument, e.g., it can be a weighted average of the action payoff and self-esteem payoff. All functions that are used are differentiable.⁴

The product structure of the state space is inessential for our analysis. One can also easily generalize the model to situations in which the agent has only one payoff index (instead of somewhat specific u and v , which are later aggregated by ϕ), and this index is a function of her action and the state of the world.

⁴This disallows, for example, u to be $-|p_t - p|$, because this function is not differentiable at $p_t = p$. However, u being $-|p_t - p|^2$ already has the required property.

However, we believe that the product structure of the state space, and the specific payoff structure facilitate thinking in terms of our model.

It may seem so far that the agent can affect only her action payoff, whereas she is born of some ability θ , and her self-esteem payoff may not be affected by anything she can do. This would indeed be true if the agent’s beliefs about the state of the world were represented by a single probability distribution over states (ω, θ) . However, this may not be the case when signals are vague. The key new feature of our model is that the agent’s action may affect her self-esteem payoff. Intuitively, we have in mind situations in which by taking some action the agent will reveal (perhaps only to herself) or admit something about her ability. So, the agent makes a conscious choice of action taking into account this effect on her self-esteem payoff.

For example, if after numerous casual observations of a child’s reading skills, an agent dramatically changes her opinion about the skills in response to learning the result of a more formal reading test, then she reveals her rather low ability of learning from casual observations, or admits some distrust to her ability of learning from casual observations; similarly, if she dramatically change her opinion about a researcher with whom she has interacted for a long time in response to learning some statistics of what the researcher has accomplished.

In terms of our model, the agent’s beliefs evolve over time as follows. At the beginning of period t , the agent is equipped with a compact convex set of probability distributions $I_{t-1} \subset \Delta(\Omega \times \Theta)$, which represents the agent’s multiple beliefs at this time. The agent first updates its beliefs about the state of the world to another compact set of probability distributions $I'_t \subset \Delta(\Omega \times \Theta)$ in response to observing the realization x_t of signal L_t . This is done by applying some updating operator, which takes as input I_{t-1} , L_t and x_t , and returns as output $I'_t \subset \Delta(\Omega \times \Theta)$. In the present paper, we will assume that the updating is performed element by element, that is, for any $p_{t-1} \in I_{t-1}$ and $l_t \in L_t$ a probability distribution $p_t \in I'_t$ is obtained by applying the Bayes rule, given that the realization x_t was observed.

Next, when action p_t is taken, the set I'_t may be further updated to a compact and convex subset I''_t of the convex hull of I'_t .⁵ This subset must have the property that action p_t maximizes her action-payoff operator U across all possible actions, given the set I''_t .⁶ (In defining the property, we assumed that t is the last period of the model. In the setting in which t is not the last period, the agent must also internalize the impact of her choice of p_t and I_t , on subsequent anticipated actions. However, we will not explore this more involved setting in the present paper, leaving it for future studies.) If multiple subsets I''_t have this property, the agent chooses one arbitrarily. The choice of p_t (and I''_t) determine the agent’s action and self-esteem payoffs as follows:

- (A) The action flow payoff is the value assigned by operator U to p_t and I''_t .
- (B) The self-esteem flow payoff is equal to the value assigned by the operator V to I''_t .

⁵The purpose of assuming convexity of all probability-distribution sets is to allow the agent for assigning various weights to posteriors or lotteries, or “averaging out” over sets of posteriors or lotteries.

⁶More precisely, U depends on the marginals on Ω of the distributions from I''_t .

We interpret that this additional updating as a response to the action p_t taken by the agent. The idea is that by taking an action the agent narrows down what she believes about the state of the world. The agent is aware of the interaction between her action and her current self-esteem payoff (and possibly subsequent payoffs). By assuming that the agent can choose any set of beliefs I_t'' in the convex hull of I_t' , which is consistent with her action, we simply allow the agent to maximize her current self-esteem payoff (and possibly subsequent payoffs), given the chosen action.

Finally, we assume that $I_t = I_t'$ or $I_t = I_t''$; in the former case, we say that that updating to I_t'' is *transient*, and in the latter case, we say it is *perpetual*. Intuitively, when $I_t = I_t'$ the agent exploits the trade-off between her action payoff and self-esteem in a single period, but this does not distort her information for a longer run; in contrast, the agent's action in any single period has a permanent effect on her information when $I_t = I_t''$.

The differences between transient and perpetual updating are typically important. However in this paper (except Section 6), they will not affect the results. To be concrete, we will focus on the perpetual updating, but the arguments remain unchanged for transient updating. In Section 6, we present two applications in which it will matter if updating is transient or perpetual. Studying perpetual updating, we will assume that $I_t = I_t''$, for $i = 1, 2, \dots$, consists of a single probability distribution, that is, by the end of each period the agent summarizes her possibly vague information about the world as a single belief. This restriction is analytically convenient, but does not affect qualitatively any of our results.

Finally, we will often restrict attention, which is with a loss of generality, to some specific structures of signals L_t . Namely, suppose that any lottery from any set L_t assigns fifty-fifty chance to both signal values, contingent on the agent being dumb, and this is independent of $\omega \in \Omega$. Therefore, every signal lottery l can be described by two numbers r_+ and r_- , which, respectively, are: (1) the probability that the lottery assigns to the realization 'true' contingent on the claim being true; and (2) the probability that the lottery assigns to the realization 'false' contingent on the claim being false.

4 Imaginary confirming and discrediting effects

In this section, we will define the two effects that were driving the analysis of Examples 1 and 2, and we characterize situations in which they are, according to our analysis, observed. We will disregard any longer-run consequences of the agent's choices; such consequences typically have a very involved form, depending on the agent's expectations regarding future decision problems. More formally, we will be exploring the two-period version of our model, in which a vague signal can be received only in period 2.

We say that an agent has no self-esteem concerns if v is constant, and we say that she has self-esteem concerns if v is strictly increasing.

Definition 1 *We say that an imaginary confirming effect is observed when after observing the realization*

of a signal: (1) an agent with no self-esteem concerns would increase (decrease) its action; and (2) an agent with self-esteem concerns increases (decreases) its action even by more.

We say that an imaginary discrediting effect is observed when after observing the realization of a signal: (1) an agent with no self-esteem concerns would increase (decrease) its action; and (2) an agent with self-esteem concerns decreases its action or increases its action by less (increases its action or decreases its action by less).

These two effects are also illustrated in Figure 2 in which the left panel illustrates the former effect, and the right panel illustrates the latter effect.




Figure 2. The beliefs at the beginning of period 2 are p_1 and q_1 .

The agent with no self-esteem concerns would update the belief about the claim to p'_2 , and the agent with self-esteem concerns updates her beliefs to p_2 and q_2 .

Some instances of signals exhibiting the imaginary self-esteem effects have been indicated in Examples 1 and 2. More generally, if a vague signal L consists only of lotteries such that $r_+ = r_-$, then the confirming effect is typically observed. We will denote the class of such signals by \mathcal{C} . And if a vague signal L consists only of lotteries such that $r_+ = 1/2$ and $r_- \geq 1/2$, then the discrediting effect is typically observed. We will denote the class of such signals by \mathcal{D} .

To see why confirming is observed for signals from \mathcal{C} , and discrediting is observed for signals from \mathcal{D} , suppose the agent enters period 2 with beliefs $p_1, q_1 \in (0, 1)$ and $p_1 \neq 1/2$; assume with no loss of generality that $p_1 > 1/2$.⁷ Suppose that first the agent obtains (in period 2) a signal consisting of a single lottery

⁷More precisely, we could alternatively assume that $p_1 < 1/2$ and study lotteries such that $r_+ \geq 1/2$ and $r_- = 1/2$.

which assigns probability r_+ to the realization ‘true’ contingent on the claim being true, and the agent being smart; and probability r_- to the realization ‘false’ contingent on the claim being false, and the agent being smart. Then, the agent’s posteriors are:

$$p_2 = \frac{p_1[q_1 r_+ + (1 - q_1)/2]}{p_1[q_1 r_+ + (1 - q_1)/2] + (1 - p_1)[q_1(1 - r_-) + (1 - q_1)/2]}$$

and

$$q_2 = \frac{q_1[p_1 r_+ + (1 - p_1)(1 - r_-)]}{q_1[p_1 r_+ + (1 - p_1)(1 - r_-)] + (1 - q_1)/2}.$$

In particular, performing some calculations, we have that

$$\frac{q_2 - q_1}{p_2 - p_1} = \frac{\left[p_1 r_+ + (1 - p_1)(1 - r_-) - \frac{1}{2} \right]}{p_1(1 - p_1) \left[\frac{r_+}{1 - q_1} - \frac{1 - r_-}{1 - q_1} \right]}. \quad (1)$$

When $r_+ = r_-$, formula (1) becomes

$$\frac{q_2 - q_1}{p_2 - p_1} = \frac{(2r_+ - 1)(2p_1 - 1)(1 - q_1)}{2p_1(1 - p_1)(2r_+ - 1)} = \frac{(2p_1 - 1)(1 - q_1)}{2p_1(1 - p_1)}. \quad (2)$$

This expression is positive, and independent of the specific value of $r_+ = r_-$.

Now, consider a vague signal L_2 composed of lotteries such that $r_+ = r_-$. The set I'_2 for this signal lies on the line passing through (p_1, q_1) with the slope given by (2). Since L_2 is vague, this set consists of more than one point. Now, our assumptions on the agent’s payoff imply that the confirming effect must be observed, unless the agent with no self-esteem concerns takes the action which corresponds to the extreme right point of I'_2 .

When $r_+ = 1/2$ and $r_- \geq 1/2$, formula (1) becomes

$$\frac{q_2 - q_1}{p_2 - p_1} = \frac{(1 - p_1)(1 - 2r_-)2(1 - q_1)}{2p_1(1 - p_1)(2r_- - 1)} = \frac{(q_1 - 1)}{p_1}. \quad (3)$$

This expression is negative, and independent of the specific value of r_- . So, the discrediting effect must be observed, unless the agent with no self-esteem concerns takes the action which corresponds to the extreme left point of I'_2 .

In general, it is relatively easy to determine which effect will be observed. We offer the following algorithm for detecting the imaginary self-esteem effects:

1. Determine the belief about the claim of the agent with no self-esteem preference. Denote this belief by p'_2 .
2. Update the prior beliefs (p_1, q_1) element by element over all lotteries belonging to the signal L_2 , and partition the posterior beliefs into those which give a lower belief about the claim than p'_2 , and those which give a higher belief about the claim than p'_2 . Denote these two sets by I_2^- and I_2^+ .
3. Take the posterior from I_2^- with the highest belief about ability, and the posterior from I_2^+ with the highest belief about ability. If there are multiple such posteriors, take any one.

4. If the so selected posterior from I_2^+ yield a higher belief about ability, then we observe an imaginary confirming effect. If the so selected posterior from I_2^- yield a higher belief about ability, then we observe an imaginary discrediting effect. If the two beliefs about ability are equal, then no effect is observed.

In Examples 1 and 2, we have computed that $(p_1, q_1) = (5/8, 1/2)$; $p'_2 = 11/18$ in Example 1 and $p'_2 = 59/84$ in Example 2. In both examples, L_2 consists of two lotteries, and updating element by element, we obtain that:

$$I_2^- = \left\{ \left(\frac{1}{2}, \frac{3}{7} \right) \right\} \text{ and } I_2^+ = \left\{ \left(\frac{13}{18}, \frac{5}{9} \right) \right\}$$

in Example 1, and

$$I_2^- = \left\{ \left(\frac{4}{7}, \frac{19}{35} \right) \right\} \text{ and } I_2^+ = \left\{ \left(\frac{5}{6}, \frac{1}{2} \right) \right\}$$

in Example 2. Since $3/7 < 5/9$, we observe an imaginary confirming effect in Example 1, and since $19/35 > 1/2$, we observe an imaginary discrediting effect in Example 2.

We will conclude this section with some comparative statics results. We will first explore the effect of extra signal vagueness. We will compare only signals L and \tilde{L} such that the agent with no preference for self-esteem takes the same action after observing a realization x generated by signal L , and after observing this realization x generated by signal \tilde{L} . We define signal L to be *more vague* than signal \tilde{L} when the convex hull of the set of probability distributions I'_2 , obtained from I_1 by updating element by element over the lotteries from L contains the convex hull of the set of probability distributions \tilde{I}'_2 , obtained from I_1 by updating element by element over the lotteries from \tilde{L} .⁸

Although signal vagueness is essential for the existence of imaginary effects studied in the present paper, more vague signals not necessarily generate stronger imaginary effects. This seems very intuitive. If we “include” to a signal a lottery which reveals good news about the agent’s ability (for example, a positive result of some ability test), then the agent may have lower incentives for “trading off” her action payoff for her self-esteem payoff.

It turns out that if this kind of additional vagueness is absent, more vague signals indeed generate stronger imaginary effects. Consider again two signals L and \tilde{L} , such that \tilde{L} is more vague than L , and a realization $x \in X$. We say that \tilde{L} contains no free good news compared to L if there is no probability distribution $(\tilde{p}, \tilde{q}) \in \tilde{I}'_2$, where \tilde{I}'_2 is obtained from I_1 when the agent observes signal \tilde{L} , and a probability $(p, q) \in I'_2$, where I'_2 is obtained from I_1 when the agent observes signal L , such that

$$p = \tilde{p} \text{ and } \tilde{q} > q.$$

⁸We could define by inclusion more vague signals in general, without restricting attention to pairs which induce the same action with no preference for self-esteem. However, for general pairs of signals, information can be more vague in a biased manner, into the positive or negative direction. Then, an agent with no self-esteem concerns typically takes different actions under different signals, and this effect interferes with our imaginary effects.

Proposition 1 *If signal \tilde{L} is more vague than signal L , and \tilde{L} contains no free good news compared to L , then the imaginary effects are (weakly) stronger under \tilde{L} than under L .*

Proof. The proof is illustrated in the left panel of Figure 3. Suppose that under L , the belief (p_2, q_2) is chosen from a segment S from the boundary of I'_2 . The ‘no free good news’ condition implies that I'_2 under signal \tilde{L} lies below the line passing through S . This in turn implies that the belief (p_2, q_2) under \tilde{L} must be either the same as under L , or be chosen not from S . In the latter case, the belief (p_2, q_2) under \tilde{L} must be placed to the right of the belief (p_2, q_2) under L if the imaginary effect was confirming, and must be placed to the left of the belief (p_2, q_2) under L if the imaginary effect was discrediting. ■




Figure 3. (left panel) The ‘no free good news’ property implies that I'_2 for signal \tilde{L} lies below the bold line passing through segment S , and to the right of (p_2, q_2) for signal L .

(right panel) By adding a lottery \tilde{l} such that posterior (\tilde{p}, \tilde{q}) is above S , we make the dashed part of the contour of I'_2 that contains (p_2, q_2) becomes flatter than S .

One may also wonder what is the impact of a more vague signal if some lottery \tilde{l} is added to signal L , such that \tilde{l} contains “free good news;” more specifically, suppose that the posterior (\tilde{p}, \tilde{q}) generated by this lottery is such that $p = \tilde{p}$ and $\tilde{q} > q$ for some (p, q) that belongs to the segment S from the boundary of I'_2 under L such that $(p_2, q_2) \in S$. This condition is, for example, always satisfied when $L \in \mathcal{C} \cup \mathcal{D}$. Assume

that the agent’s flow payoff is a concave function of q , which seems to be the most realistic case. And to avoid considering too numerous cases, suppose that the confirming effect was observed under L .

Suppose first that $\tilde{p} < p'_2$ and $\tilde{q} < q'$ some other $(p', q') \in I'_2$. Recall that p'_2 is the action which would be chosen by an agent with no self-esteem concerns. Such a case is illustrated in the right panel of Figure 3, and intuitively, such a lottery \tilde{l} contains moderate good news about ability contingent on interpreting the signal as negative. In this case, the confirming effect is (strictly) weaker under \tilde{L} than under L . Indeed, adding lottery \tilde{l} implies that the agent chooses (p_2, q_2) from a segment that is flatter than, and lies above the segment from which she has chosen (p_2, q_2) under signal L . This, combined with the assumption that indifference curves become steeper when q increases, yield the desired result. The result is illustrated in the right panel of Figure 3. Intuitively, the agent must sacrifice more in terms of her action payoff to obtain a higher self-esteem payoff, and in addition, she values an extra self-esteem payoff less, since the marginal value of this payoff diminishes with q .

If $\tilde{p} < p'_2$ and $\tilde{q} \geq q'$ for all other $(p', q') \in I'_2$, which intuitively means that the additional lottery contains very strong good news about ability contingent on interpreting the signal as negative, the confirming effect disappears under \tilde{L} , or even the discrediting effect is observed. Finally, if $\tilde{p} \geq p'_2$, which intuitively means that good news about ability come with interpreting the signal as positive, then the impact on the imaginary confirming effect is ambiguous. Intuitively, the ambiguity comes from two opposite effects. The agent must sacrifice less in terms of her action payoff to obtain a higher self-esteem payoff, but she values an extra self-esteem payoff less, since the marginal value of this payoff diminishes with q .

We will now explore the effect of the agent’s prior beliefs. We will consider only signals from classes \mathcal{C} and \mathcal{D} . The results can be generalized to other signals, by referring to our algorithm, but these more general results would require additional assumptions, and would be less transparent.⁹ Assume with no loss of generality that $p_1 \geq 1/2$.

Proposition 2 (i) *Suppose that the agent’s self-esteem payoff is a concave function of q . Then an increase in q_1 reduces the imaginary confirming (discrediting) effect, given any signal $L_2 \in \mathcal{C}$ ($L_2 \in \mathcal{D}$).*

(ii) *An increase in p_1 magnifies the imaginary confirming effect and reduces imaginary discrediting effect, given any signal $L_2 \in \mathcal{C}$ and $L_2 \in \mathcal{D}$, respectively.*

Proof. Consider first the confirming effect, and the lotteries such that $r_+ = r_-$. Expression (2) increases in p_1 , and decreases in q_1 . This combined with the fact that indifference curve become steeper when q_1 increases, and the shape of indifference curves remains unaltered when p_1 increases yield the desired result (see the left panel of Figure 4).

⁹The problem with obtaining more general results is again that a change in priors typically “translates” the set I'_2 , and this may change the type of imaginary effect.

For signals from classes \mathcal{C} and \mathcal{D} , if the belief (p_2, q_2) belongs to some segment from the boundary of I'_2 , then the belief (p_2, q_2) after the prior change belongs to the “image” of this segment in the translated I'_2 .

If this last property is assumed, then our proposition generalizes to other signals.

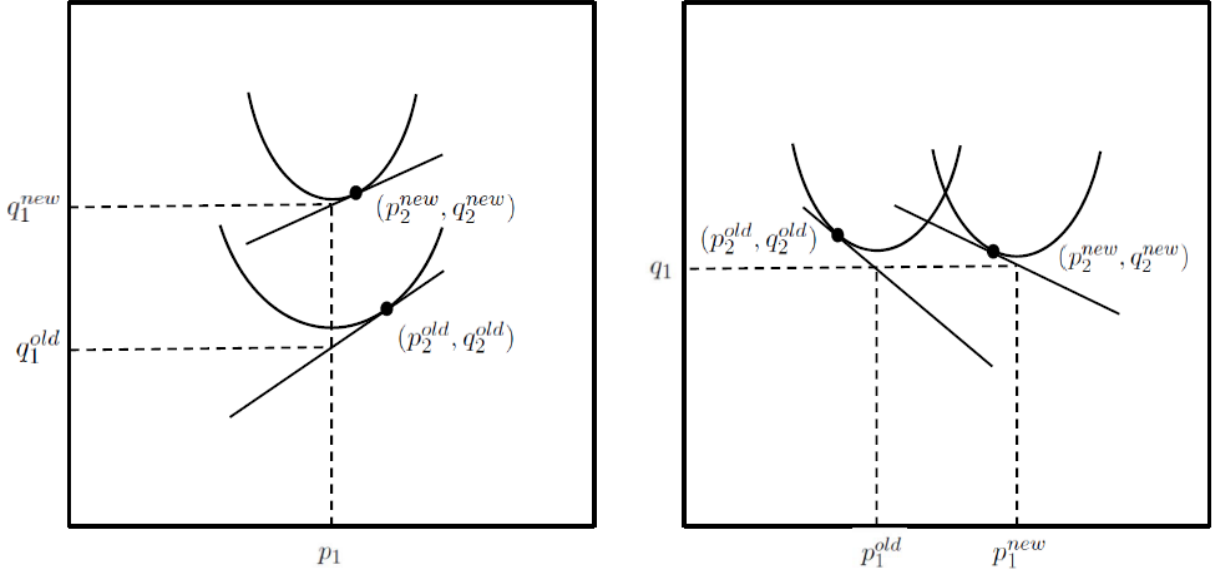


Figure 4. The effect of an increase in q_1 on the imaginary confirming effect is illustrated in the left panel, and the effect of an increase in p_1 on the imaginary discrediting effect is illustrated in the right panel.

Consider now the discrediting effect, and the lotteries such that $r_+ = 1/2$ and $r_- \geq 1/2$. Expression (3) decreases in absolute value in p_1 , and decreases in absolute value in q_1 . This combined with the fact that indifference curves become steeper when q_1 increases, and the shape of indifference curves remains unaltered when p_1 increases yield the desired result (see the right panel of Figure 4). ■

5 Temporary or permanent effects

In the standard model in which agents have a single belief over the states of the world, when agents obtain an infinite sequence of informative signals, and the information contained by individual signals is uniformly bounded away from zero, agents gradually learn the state of the world. One may wonder whether agents will gradually learn the state of the world in our multi-belief model, after receiving a sequence of informative signals containing information that is uniformly bounded away from zero.

It is also interesting to know if our model predicts polarization as a short-lived phenomenon, or rather disagreement is a “limit” phenomenon, that is, the beliefs of agents may be diverging even when they are

observing an infinite sequence of public informative signals. Finally, Benoit and Dubra (2011) cite various evidence from interviewing professional truck drivers, and according to this evidence, younger drivers rate themselves as being better overall drivers than their peers, while older drivers rank themselves comparably to drivers in their age group. So, it may be interesting to see whether this kind of evidence is consistent with our analysis.

We provide an example showing that agents may never learn the state of the world, and disagreement may happen to be a permanent phenomenon, even when signals contain a “grain of truth” (that is, information which is, in our interpretation, uniformly bounded away from being negligible), but we also provide a result showing that agents learn the true state of the world if every signal contains a grain of truth, but this grain of truth must, in our interpretation, be quite large.

Example 1 (continued). Suppose that the signals are exactly as in Example 1. Say that the realizations of the first signal was ‘true’. Represent now the second signal as a set of two lotteries:

- lottery 1 assigning probability $3/4$ to the logical value of the claim, contingent on the agent being smart, and probability $1/2$ to each value, contingent on the agent being dumb;
- lottery 2 assigning probability $3/4$ to the opposite value (opposite to the logical value of the claim), contingent on the agent being smart, and probability $1/2$ to each value, contingent on the agent being dumb.

Note that the set I_1 (of agent T) consists of the single joint probability distribution π_1 which assigns probabilities:

$$\left[\begin{array}{l} \frac{3}{8} \text{ to true, smart;} \\ \frac{2}{8} \text{ to true, dumb;} \\ \frac{1}{8} \text{ to false, smart;} \\ \frac{2}{8} \text{ to false, dumb.} \end{array} \right]$$

In dependent of the realization of the second signal, the set I'_2 , assuming element by element updating, consists of two joint probability distributions π'_2 and π''_2 , which assign probabilities:

$$\left[\begin{array}{l} \frac{63}{126} \text{ to true, smart;} \\ \frac{28}{126} \text{ to true, dumb;} \\ \frac{7}{126} \text{ to false, smart;} \\ \frac{28}{126} \text{ to false, dumb,} \end{array} \right]$$

and

$$\left[\begin{array}{l} \frac{27}{126} \text{ to true, smart;} \\ \frac{36}{126} \text{ to true, dumb;} \\ \frac{27}{126} \text{ to false, smart;} \\ \frac{36}{126} \text{ to false, dumb,} \end{array} \right]$$

respectively. The convex combination of these two probability distributions with coefficients $9/16$ and $7/16$ is equal to the probability distribution π_1 . So, the set I_2 may consist of the single probability distribution π_1 . If the signal L_t in every period $t = 3, 4, \dots$ is the same as that in period 2, then I_t will consist of the single probability distribution π_1 for all $t = 1, 2, \dots$. This implies that the agent will never learn the logical value of the claim, in particular, disagreement of agents T and F will be a permanent phenomenon.

The signal in period 2 contains no grain of truth. Indeed, there is no sense in which one could claim that it makes one of the two states appear more likely. A grain of truth can, however, be added by continuity. Indeed, for a small $\varepsilon > 0$ replace the signal L_2 from Example 1 with an alternative pair of lotteries:

- lottery 1 assigning probability $3/4$ to the logical value of the claim, contingent on the agent being smart, and probability $1/2$ to each value, contingent on the agent being dumb;
- lottery 2 assigning probability $3/4 - \varepsilon$ to the opposite value (opposite to the logical value of the claim), contingent on the agent being smart, and probability $1/2$ to each value, contingent on the agent being dumb.

Next, suppose that the claim is false, and agent T obtains signal ‘false’ in period 2. The posterior π_2'' remains the same, while the posterior π_2' is slightly changed. However, it is still the case that a convex combination of π_2' and π_2'' yields π_1 . The coefficients must only be slightly higher than $9/16$ and slightly lower than $7/16$, respectively.

Thus, agent T observing the realization ‘false’ in periods $t = 2, 3, \dots$ maintains the estimate of the claim being true at $5/8$, and obviously the agent even raises the estimate when some realizations are ‘true.’ This, and analogous arguments for agent F, imply that no agent will ever learn the logical value of the claim, and the two agents will forever disagree.

Note finally that the agent will not learn her actual ability either.

For the result, assume that every signal lottery assigns the same probability to realization ‘true,’ when the value of ω is true, and to realization ‘false,’ when the value of ω is false. That is, we will consider signals from class \mathcal{C} . (An analogous result, by an analogous argument, holds for signals from class \mathcal{D} .) Then, for any signal $L \subset \Delta(X \times \Omega \times \Theta)$ define r_H as the highest r over all lotteries from the set, and r_L as the lowest r over all lotteries from the set, that is,

$$r_H = \max\{r_l : l \in L\} \text{ and } r_L = \min\{r_l : l \in L\}.$$

We can with no loss of generality assume that $1 - r_H \leq r_L$; otherwise, we could rename signal value ‘false’ for ‘true,’ and signal value ‘true’ for ‘false,’ without affecting the information that agents obtain. That is, the lottery with r_H is most informative in L ; when $r_L \geq 1/2$, the lottery with r_L is least informative, but when $r_L < 1/2$, the lottery with r_L may not be least informative in L , but transmits information in a perverse way: the value ‘false’ suggests that the claim is true, and the value ‘true’ suggests that the claim is false.

Proposition 3 *Suppose that $r_L^t < 1/2$ only for a finite number of signals L_t , and for some $\eta > 0$, $1/2 + \eta < r_L^t$ for an infinite number of signals L_t . Then, for any true ω , and any probability distributions over sets L_t , $t = 1, 2, \dots$, according to which Nature chooses signal lotteries, any agent’s belief about ω converges with probability 1 to the true ω .*

It obviously follows from that under the assumptions of the proposition the agent’s beliefs about θ converges to the actual value of θ , because the agent who knows ω , and the structure of signals, will eventually learn θ .

Proof. Since the prior has full support, the agent still assigns a positive probability to each state of the world after observing a finite number of signals L_t with $r_L^t < 1/2$. After observing all signals L_t with $r_L^t < 1/2$, the agent observes an infinite sequence of signals L_t with $1/2 + \eta < r_L^t$. All lotteries of these signals contain information, and this information is bounded away from 0. (After observing all signals L_t with $r_L^t < 1/2$, the agent may also observe signals with $r_L^t = 1/2$, which contain lotteries with no information.) Thus, by the same argument as in the standard model with a single belief, the agent will update beliefs, so that they will converge in expectation to the state of the world. Moreover, by the Hoeffding inequality, these beliefs will uniformly converge to the state of the world with probability 1, across all probability distributions over sets L_t , $t = 1, 2, \dots$, according to which Nature chooses signal lotteries. ■

In (the continuation of) Example 1, $r_L^t = 1/4$ (or $1/4 + \varepsilon$) and $r_H^t = 3/4$ for $t = 2, 3, \dots$, so the assumption that $r_L^t < 1/2$ only for a finite number of signals L_t is violated, and Proposition 1 does not hold in this case. The purpose of both our example and proposition is to illustrate some phenomena, not to provide any general (necessary and sufficient) conditions for the phenomena. One may wonder, however, what such necessary and sufficient conditions would be. The conditions assumed in our proposition can obviously be somewhat weakened. However, we believe that the necessary and sufficient conditions are likely to be involved, and finding them would rather be an exercise in series theory than in economics. So, we will leave these conditions unexplored, believing that the present analysis gives a good idea what is the “grain of truth” that the signals must contain in order to guarantee that the agent will over time learn the state of the world.

6 Other applications

In this section, we provide some additional examples of how various behavioral evidence may be interpreted in our model. These examples will also illustrate differences between transient and perpetual updating.

Example 3. Example 1 already suggests an argument for confirmatory bias, discussed in Rabin and Schrag. The second signal there was neutral, should be interpreted as containing no information. Yet, agent T increased the probability assigned to the claim being true. Moreover, we modified Example 1 in Section 5 so that it seemingly delivers negative information about the claim, but agent T maintains the probability assigned to the claim being true.

We will now describe a stronger instance of confirmatory bias, a situation in which the agent receives two signals, and her assessment of the correctness of the claim, whether it is more likely to be true or false, depends on the ordering in which the two signals were received. So, the example shows the role of first impression in its exact (or pure) form.

As in Example 1, suppose that an agent has fifty-fifty independent priors about the claim and her ability. Unlike in Example 1, consider a single agent, and assume that both signals (one received in period 1 and one received in period 2) are vague. More specifically, suppose that after receiving the signal ‘true’ in period 1 the agent cannot tell whether its value should be interpreted according to the structure described in Example 1, or the signal is meaningless. Recall that the former interpretation means that if the agent is smart, the value of a signal is equal to the logical value of the claim with probability $3/4$; if the agent is dumb, the value is equal to the logical value of the claim only with probability $1/2$. The latter interpretation means a fifty-fifty chance of taking each value, independent of the state of the world. In turn, after receiving the signal ‘false’ in period 2 the agent cannot tell whether its value should be interpreted according to the structure described in Example 1, or the signal is meaningless.¹⁰

Consider first transient updating. Then, the first impression does not matter. Indeed, independent of the ordering of signals, the set I_2 is the convex hull of four probability distributions: $(1/4, 1/4, 1/4, 1/4)$, $(3/8, 1/4, 1/8, 1/4)$, $(1/8, 1/4, 3/8, 1/4)$ and $(3/14, 4/14, 3/14, 4/14)$, where as always the first coordinate refers to state (true, smart), and the remaining coordinates to states (true, dumb), (false, smart) and (false, dumb). The first probability distribution was obtained assuming that both signals were meaningless; the second and third ones were obtained assuming that one of the signals was informative and the other meaningless, and the fourth one was obtained assuming that both signals were informative. For the rest of the analysis, assume perpetual updating.

Suppose for a moment that the agent obtains utility only from making the right assessment of the claim, and whenever she has multiple beliefs she takes the average of the minimum and maximum assessment. Suppose that the agent transmits from period 1 to period 2 the single probability distribution $(5/16, 1/4, 3/16, 1/4)$ obtained by taking the average of the two beliefs which the agent has after observing the first signal: $(3/8, 1/4, 1/8, 1/4)$ and $(1/4, 1/4, 1/4, 1/4)$. After observing the second signal, the agent forms two beliefs: $(5/16, 1/4, 3/16, 1/4)$ and $(5/30, 8/30, 9/30, 8/30)$. Thus, the new assessment of the claim being true is the average of $9/16$ and $13/30$, and this is equal to $239/480 < 1/2$, and so no confirmatory bias is observed. Actually, the agent decreases the probability assigned to the claim being true below $1/2$.¹¹ The

¹⁰These signals can be represented as sets of two lotteries: one taking the logical value of the claim with probability $3/4$ when the agent is smart, and taking the logical value of the claim with probability $1/2$ when the agent is dumb, and the other taking each value with probability $1/2$, independent of the state of the world.

¹¹This feature of the model follows from the fact that Bayesian updating downwards is a convex operation, and Bayesian updating upwards is a concave operation, while taking the average over two beliefs is a linear operation. In particular, if the order of observing the two signals was reversed, the agent would increase the probability assigned to the claim being true to above $1/2$.

new assessment of being smart is $232/480$.

Suppose now that the agent has in addition a preference for being assessed as smart, then the pair of second-period assessments moves along the segment joining $(9/16, 1/2)$ and $(13/30, 7/15)$ towards pair $(9/16, 1/2)$, showing some confirmatory bias. The exact position of the new pair of assessments on the segment depends on the marginal rate of substitution between the preference for assessing the claim correctly and being assessed as smart, even reaching pair $(9/16, 1/2)$ in an extreme. So, the agent may even simply ignore the evidence against the initial impression, exactly as is assumed in Rabin and Shrag.

Of course, we would obtain the opposite conclusions if the ordering of signals was reversed, reaching $(7/16, 1/2)$ in an extreme, and showing that the first impression matters. What is the intuition behind the role of the first impression and the difference between transient and perpetual updating? Under transient updating, the agent may have biased beliefs in period 1, but she does not carry the bias forward to period 2. This enables her to realize that the signals are symmetric, and no opinion about the claim (positive or negative) favors the view that her ability is high. Under perpetual updating, the agent carries her biased beliefs to the second period, and maintains the beliefs biased in the same directions as those from the first period, trying to maintain the view that her ability is high.

Example 4. In this example, we show that agents with more extreme views adopt more extreme opposite views when they decide to change their views, or that overconfident agents change their views more substantially. There seems to be quite rich historic and anecdotal evidence that this indeed takes place in practice. However, the only available empirical evidence seems to be that on the relation between overconfidence and political extremism contained in Ortoleva and Snowberg (2015).

Consider the signal structure from Example 1, that is, each signal takes one of two possible values: ‘true’ or ‘false.’ If the agent is smart, the value of a signal is equal to the logical value of the discussed claim with probability $3/4$; if the agent is dumb, the value is equal to the logical value of the discussed claim only with probability $1/2$. Unlike in Example 1, suppose that the agent enters period 1 with priors $5/8$ and $1/2$. (In Example 1, these were the agents’ beliefs after period 1. So, one can think that the agent observed in period 0 the signal that in Example 1 she observed in period 1.)

Suppose that in period 1 the agent observes the vague signal that in Example 1 she observed in period 2. That is, the signal has the same structure as the signal in period 1, but the agents cannot tell what was the signal’s value, whether it was ‘true’ or it was ‘false,’ even they cannot tell what is the chance of the signal taking each of the two values. By Example 1, the agent chooses the assessment $p_1 = p$ of the claim being true with from interval $[1/2, 13/18]$; this induces her assessment $q_1 = q$ of being smart as $q = 4p/7 + 1/7$, and higher is the agent’s relative payoff of being assessed as smart versus her action payoff, the higher p_1 and q_1 she picks.

We suggest not to assign to this finding any meaning, because the averaging over beliefs is rather an ad-hoc rule. One could eliminate this kind of effects by assuming a little more complicated decision rule under multiple priors, which would take into account the convexities and concavities of Bayesian updating.

Finally, suppose that in period 2 the agent observes two “candid” independent signals which have the same structure as the signal from period 0. (If one wishes to allow for only one signal per period, as we have done till now, suppose that the two signals are observed in periods 2 and 3.) Suppose that the value of both signals was ‘false’.

Under transient updating, higher was the p_1 picked in period 1, lower is the p_2 at the end of period 2. In particular, a higher is relative payoff of the agent from being assessed as smart versus her action payoff, more extreme are the changes in the agent’s views. Indeed, the choice of p_1 in period 1 is based on one positive candid signal, and one extremely vague signal. Similarly, the choice of p_2 in period 2 is based on one positive candid signal, two negative candid signals, and and one extremely vague signal. Since one positive signal and one negative signal “cancel out,” the choice in period 2 is a mirror image of that in period 1, which proves the claim.

The conclusion is ambiguous under perpetual updating. Indeed, a higher $p_1 = p$ chosen in period 1 implies a higher $q_1 = q$, as $q = 4p/7 + 1/7$. The higher p has a positive effect on the assessment of the claim in period 2, as the two negative signals must be combined with the more positive belief inherited from period 1. However, the higher q has a negative effect on the assessment of the claim in period 2, as the negative signals are, according to the agent, more informative. For the parameter values studied in Example 1, and so in the present example, one can compute that a higher p_1 chosen in period 1 implies a higher p_2 in period 2 (the former being greater than 0, and the latter being smaller than 0).

Numerous other phenomena that resemble what we observe in practice can be generated in the present setting. For example, Piotr Dworzak reported to the author that a referee who has a preference for being viewed as smart may write a positive report despite disliking the paper under review. The reason is that the referee may interpret a negative but vague signal obtained while reading of the paper as an evidence of its low quality or an evidence of low writing skills of the author, and this latter view may “reveal” better news about the referee as being smart.

7 Information Acquisition

Throughout the paper, we were assuming exogenous information structure. The model allows, however, for studying some form of information acquisition, or some choice of signals. This is, in particular, an important advantage of our model over that from Brunnermeier and Parker (2015). However, allowing for endogenous signal structure raises some conceptual issues. Agents in our model always prefer vague signals. So, they could in principle purposely choose more vague signals to raise their utility.

This is, however, unclear to what extent real-life agents would feel smarter only because they have chosen more vague signals, which they could have later interpreted as good news about their ability. Rather, they would at least partially filter out the additional vagueness coming from their choice. In addition, some agents may exhibit an aversion to making decisions under vague information. We leave the issue of information

acquisition for future research.

8 References

Acemoglu, D., V. Chernozhukov and M. Yildiz. (2009) “Fragility of Asymptotic Agreement under Bayesian Learning,” *Theoretical Economics*, forthcoming.

Ahn D. (2008) “Ambiguity Without a State Space,” *Review of Economic Studies* 75(1): 3-28.

Andreoni, J. and T. Mylovannov (2012) “Diverging Opinions,” *American Economic Journal: Microeconomics* 4: 209–32.

Arrow K.J. and Hurwicz L. (1972) “An Optimality Criterion for Decision-Making under Ignorance,” in C.F.Carter and J.L.Ford (Eds.), *Uncertainty and Expectations in Economics: Essays in Honour of G.L.S. Shackle*, Basil Blackwell, Oxford, pp. 1-11.

Baliga, S., E. Hanany and P. Klibanoff (2013) “Polarization and Ambiguity,” *American Economic Review* 103(7): 3071–3083.

Bénabou R. and Tirole J. (2002) “Self-Confidence and Personal Motivation,” *Quarterly Journal of Economics* 117(3): 871–915.

Benoit J.P. and J. Dubra (2011) “Apparent Overconfidence,” *Econometrica* 79 (5): 1591–1625.

Brunnermeier, M. K. and J. A. Parker (2005) “Optimal Expectations,” *American Economic Review* 95 (2005): 1092-1118.

Darley, J. M. and P. H. Gross (1983) “A Hypothesis-Confirming Bias in Labeling Effects,” *Journal of Personality and Social Psychology* 44: 20–33.

Dunning, D. (1989) “Ambiguity and Self-Evaluation: The Role of Idiosyncratic Trait Definitions in Self-Serving Assessments of Ability,” *Journal of Personality and Social Psychology* 57(6): 1082-1090.

Greenwald, A.G. (1980) “The totalitarian ego: Fabrication and revision of personal history,” *American Psychologist* 35: 603-618.

Kondor, P. (2012) “The More We Know about the Fundamentals, the Less We Agree on the Price,” *Review of Economic Studies* 79: 1175–1207.

Olszewski W. (2007) “Preferences over Sets of Lotteries,” *Review of Economic Studies* 74 (2007): 567-595.

Ortoleva P. and E. Snowberg (2015) “Overconfidence in Political Behavior,” *American Economic Review* 105(2): 504-535.

Rabin, M. and J. Schrag (1999) “First Impressions Matter: A Model of Confirmatory Bias,” *Quarterly Journal of Economics* 114(1): 37-82.

Santos-Pinto, L. and J. Sobel, (2005) “A Model of Positive Self-Image in Subjective Assessments,” *American Economic Review*, 95(5): 1386–1402.

Svenson, O. (1981) “Are we all less risky and more skillful than our fellow drivers?” *Acta Psychologica* 47(2): 143–148.

- Van den Steen, E. (2004) “Rational overoptimism,” *American Economic Review*, 94(4): 1141–1151.
- Vierø M-L. (2009) “Exactly What Happens After the Anscombe—Aumann Race? - Representing Preferences in Vague Environments,” *Economic Theory* 41: 175-212.
- Zábojník, J. (2004) “A model of rational bias in self-assessments,” *Economic Theory* 23(2): 259-282.
- Zimper, A. and A. Ludwig. (2009) “On Attitude Polarization under Bayesian Learning with Non-additive Beliefs,” *Journal of Risk and Uncertainty* 39: 181–212.