WILL YOU NEVER LEARN? SELF DECEPTION AND BIASES IN INFORMATION PROCESSING[1]

DANIEL GOTTLIEB

Wharton School, University of Pennsylvania
dgott@wharton.upenn.edu

Abstract

This paper considers a repeated model of selective awareness and studies implications for information processing. An individual receives a sequence of signals. Each signal is informative about the "state of the world." As in the static models of Benabou and Tirole (2002, 2004) and Benabou (2008a, 2008b), the individual can choose how to interpret each signal. The individual's behavior displays patterns consistent with observed biases in information processing. She displays a tendency to interpret information in ways that support original beliefs and attaches a disproportionately large weight to initial observations (confirmation bias). She also updates beliefs in the right direction, but in insufficient amount compared to the update derived by Bayes' rule (conservatism bias). Additionally, the individual disregards information after a certain number of observations. As a consequence, learning is always incomplete.

1 Introduction

Economists typically model humans as statisticians who collect information in an unbiased manner and make impartial inferences. The psychological evidence, however, suggests that we tend to behave like unscrupulous statisticians, who collect and interpret information interested more in feeling competent than in the accuracy of our inferences. Sedikides, Green, and Pinter (2004, pp. 165), for example, describe people as "striving for a positive self-definition or the avoidance of a negative self-definition (...) at the expense of accuracy and truthfulness."

Social scientists have long recognized that our biases in information processing may lead to imperfect learning. Montier (2007), for example, argues that "the major reason we don't learn from our mistakes (...) is that we simply don't recognize them as such. We have a gamut of mental devices all set up to protect us from the terrible truth that we regularly make mistakes."

[1] I thank Dilip Abreu, Roland Benabou, Stephen Morris, John Smith, and Muhamet Yildiz for valuable suggestions, and seminar participants at Princeton, London School of Economics, the Behavioral Economics Conference (Cornell), and the Murray S. Johnson Memorial Conference (UT Austin) for comments.

This version: May, 2010.

This paper studies how biases in information processing arise and persist when individuals are subject to selective awareness (either by selective attention or selective memory). In order to analyze this issue, I develop a model based on the framework of Benabou and Tirole (2002, 2004, 2006a, 2006b) and Benabou (2008a, 2008b). The model provides a unified explanation for certain biases in information processing.

Consistently with a huge amount of psychological evidence, the model predicts a tendency to seek and interpret information in ways that support original beliefs (confirmation bias). According to Evans (1989, pp. 41), "[c]onfirmation bias is perhaps the best known and most widely accepted notion of inferential error to come out of the literature on human reasoning." As a result, individuals in my model attach a disproportionately large weight to initial observations. Individuals with a favorable initial streak of information become trapped with optimistic beliefs and disregard negative information received afterwards. Similarly, individuals with a negative initial streak of information are stuck into a pessimistic trap, in which every positive information is disregarded. The model also leads to an updating bias known as conservatism. Conservatism states that individuals update their beliefs in the right direction, but by too little relative to the update derived by Bayes' rule.

The model studies an infinitely lived individual, who receives a sequence of binary signals. Each signal provides information about the state of the world. A state of the world may correspond to the individual's (unknown) skills (Section 3), or some other unknown feature that affects her payoff through anticipatory utility (Subsection 4.2).

I consider two different information structures. In the first one, there is a small probability that the process will terminate in each period and the individual will have to take an action (Section 3). As Wilson (2005) argues, this setup captures an environment in which the individual expects to obtain a long sequence of information but is unsure about when the decision will have to be made. In the second information structure, the individual observes a fixed number of signals and takes an action after the last observation (Subsection 4.3). This setup represents an environment in which the individual is certain about when the decision will have to be made and how much information she will acquire before then. The payoff from the action depends on the state of the world.

The individual processes information as in the models of Benabou and Tirole. After observing each signal, she chooses whether to interpret it realistically or to rationalize it away. The main feature of the model is a trade-off between optimism and improved decision-making. When the action space is finite, the individual becomes increasingly convinced about which action to take as she observes more signals. Therefore, she chooses to rationalize negative information away after becoming sufficiently convinced of which action to take since the chance that each additional signal would affect her choice becomes arbitrarily small. When the action space is continuous, each individual signal may always affect the decision. Nevertheless, because the cost of distorting an action close to the optimum is of second order, the individual will always choose to rationalize negative information away after a sufficiently large number of observations.

Therefore, the model predicts that first impressions matter: initial information gets a disproportionately large weight. The model also predicts a confirmatory bias: individuals do not change their beliefs after they become sufficiently convinced of which action to take. Moreover, since all information is rationalized away after a certain number of observations, individuals never learn the true state. Thus, as in the quote by Montier, selective awareness creates a barrier to learning.

This paper is related to three separate literatures. The first one studies selective awareness, cognitive dissonance, self-deception, and other forms of belief manipulation. The selective awareness framework used throughout the paper has been used to provide theories of personal motivation (Benabou and Tirole, 2002), redistributive policies (Benabou and Tirole, 2006), groupthink and ideologies (Benabou, 2008a, 2008b), endowment, sunk cost effects, and some other deviations from expected utility theory (Gottlieb, 2009), and preferences for increasing payments (Smith, 2009a and 2009b).[2]

It is often argued that the Bayesian updating assumption embedded in this framework combined with the repeated nature of the decisions being modeled would lead individuals to eventually learn the truth and, therefore, the departures from rationality would vanish in the long-run. The present paper studies this argument formally. Because all information is disregarded after a certain number of observations, learning is always incomplete, and the departures from rationality presented in the static models in the literature do not disappear even when the decision problem is repeated infinitely many times.[3]

The second literature studies biases in information processing.[4] The paper closest to mine is Wilson (2005), which considers a model featuring the same information structure as the one in Section 3 and also leads to biases in information processing. The main difference between our models lies in the way memory is modeled. Wilson considers an unbiased memory, consisting of a finite number of states. This restriction precludes the individual from conditioning the action on the whole sequence of signals (or any sufficient statistic), which makes it impossible for the true state to eventually be learned. By contrast, in the model in this paper, the individual could, in principle, condition actions on the whole history of signals. The choice of not doing so arises endogenously through either the desire to improve one's self-image or to enjoy anticipatory utility. Although it is hard to dispute that human memory is bounded, it is not clear why unbiased decision makers would not be able to keep a record of their observations (say, by writing them or some sufficient statistic down) or search for evidence if needed. By contrast, the individuals considered in this paper would write down inaccurate observations, interpret them incorrectly, delete their records, or choose not to look at them.

The model can be alternatively interpreted as a costless signaling game between an expert

[2]Bernheim and Thomadsen (2005) use a similar model to show why individuals may cooperate in a prisoner's dilemma game. Kopczuk and Slemrod (2005) consider how the desire to avoid thinking about death may explain puzzles in health and savings behavior. Other papers featuring belief manipulation include Akerlof and Dickens (1982), Schelling (1985), Kuran (1993), Rabin (1994), Carrillo and Mariotti (2000), Bodner and Prelec (2002), and Di Tella et al. (2007).

[3]This paper is also related to a literature that studies learning by Bayesian decision makers. Ali (2009), for example, studies time-inconsistent individuals who face a sequence of temptations and identifies necessary and sufficient conditions for beliefs to converge to the truth. In his model, failure of learning may emerge because individuals who suspect to be time-inconsistent may prefer to commit to a certain decision. But, in this case, they never observe whether they would have resisted temptation and do not observe new information. In the present paper, the individuals always receive new information. However, incomplete learning occurs because the information gets disregarded. Acemoglu, Chernozhukov, and Yildiz (2009) obtain lack of convergence in asymptotic beliefs when individuals are uncertain about the distribution of signals.

[4]Most models in this literature focus on identifying certain biases and exploring their implications for the economic models. See, e.g., Rabin and Schrag (1999), Gennaioli and Shleifer (2008), Madarász (2009), and Schwartzstein (2009). Alternatively, Brocas and Carrillo (2009) propose a neuroeconomic model that also leads to biases in information processing.

who acquires new information in each period and has reputation concerns and an uninformed individual who has to make a decision. Therefore, the third literature to which this paper is related studies information transmission and reputation in repeated environments.[5]

The structure of the paper is as follows. Section 2 briefly reviews the psychological evidence. Section 3 introduces and discusses the main framework. In Subsection 3.1, I describe the results when actions are binary. Subsection 3.2 considers a continuum of actions. Section 4 provides generalizations and extensions of the model: Subsection 4.1 considers general prior distributions and signal structures, Subsection 4.2 presents a version of the model based on anticipatory utility, and Section 4.3 considers the case in which the individual observes a fixed number of signals. In Section 5, I discuss the reinterpretation of the model in terms of an information transmission game and its implications for information disclosure in repeated environments. Then, Section 6 concludes.

2 An Overview of the Psychology Literature

Psychologists have documented several systematic biases in how we update beliefs after receiving new information.[6] Several studies have documented a tendency to seek and interpret information in ways that support previously held beliefs. In the words of Oswald and Grosjean (2004, pp. 79):

> "Confirmation bias" means that information is searched for, interpreted, and remembered in such a way that systematically impedes the possibility that the hypothesis could be rejected – that is, it fosters the immunity of the hypothesis.

Psychologists have devoted an immense amount of work to study the confirmation bias. Indeed, according to Nickerson (1998), "[i]f one were to attempt to identify a single problematic aspect of human reasoning that deserves attention above all others, the confirmation bias would have to be among the candidates for consideration." This research suggests that as people become more convinced of their hypotheses, they tend to disregard information that conflicts with them. As Lord, Ross, and Lepper (1979, pp. 2099) summarize,

> There is considerable evidence that people tend to interpret subsequent evidence so as to maintain their initial beliefs. The biased assimilation processes underlying this effect may include a propensity to remember the strengths of confirming evidence but the weaknesses of disconfirming evidence, to judge confirming evidence as relevant and reliable but disconfirming evidence as irrelevant and unreliable, and to accept confirming evidence at face value while scrutinizing disconfirming evidence hypercritically.

In this paper, it will be useful to distinguish between two forms of confirmatory bias. I will say that individuals exhibit a *weak confirmation bias* if there is positive probability of reaching histories in which every additional information gets disregarded. I will say that individuals exhibit a *strong confirmation bias* if histories in which they disregard every additional information are reached with probability one.

[5] Papers considering reputation in games of information transmission include Sobel (1985), Benabou and Laroque (1992), Morris (2001), Morgan and Stocken (2003), and Ottaviani and Sørensen (2006a, 2006b).

[6] See Rabin and Schrag (1999) for a similar description of this literature.

The main results of this paper show that individuals with selective awareness always display confirmation biases. When actions are binary, Theorem 1 establishes the weak confirmation bias for any equilibrium and Theorem 2 establishes the strong confirmation bias for any Markovian equilibrium under uniform priors. For continuous actions, Theorem 3 shows the strong confirmation bias for any equilibrium under uniform priors. Theorems 4 and 5 generalize the strong confirmation bias results for any regular prior distribution. Theorem 6 and Proposition 5 establish the weak confirmation bias for any equilibrium in the cases of anticipatory utility and a fixed terminal period, respectively.

A related finding is the disproportionate effect of first impressions. When individuals observe sequences of exchangeable information, initial observations are excessively weighted. The Markovian equilibria of the model feature exactly this feature (see Corollaries 1 and 2 for uniform prior distributions and Corollaries 3 and 4 for any regular prior distribution). Individuals interpret initial signals realistically and update their beliefs according to Bayes' rule. However, after a certain number of periods, they discard every additional information and therefore do not update beliefs.

Since the 1960s, several psychologists have identified another updating bias known as conservatism. Conservatism states that individuals update their beliefs in the right direction but in a smaller magnitude than implied by Bayes' rule. Edwards (1968) summarizes the findings as follows:[7]

> An abundance of research has shown that human beings are conservative processors of fallible information. Such experiments compare human behavior with the outputs of Bayes's theorem, the formally optimal rule about how opinions (...) should be revised on the basis of new information. It turns out that opinion change is very orderly, and usually proportional to numbers calculated from Bayes's theorem – but it is insufficient in amount.

In Proposition 2, I show that individuals in the selective awareness model update information precisely as described by Edwards.

A large literature in psychology finds that preferences affect beliefs.[8] For example, biased beliefs seem to be more prevalent for traits and behaviors that individuals regard as important (e.g. MacDonald and Ross, 1999, Sanbonmatsu et al., 1987). Bahrick, Hall, and Berger (1996), for example, study distortions in college students's memory for their high school grades. They find that accuracy of recollections declines monotonically with the students' letter grades (from 89% for grades of A to 29% for grades of D). As Bahrick, Hall, and Berger interpret their findings, "[d]istortions are attributed to reconstructions in a positive, emotionally gratifying direction." Similarly, Greene (1981) argues that a desire to see oneself as someone who makes correct judgements biases beliefs towards one's previous decisions. Accordingly, Taylor and Gollwitzer (1995) find that biases in beliefs are stronger for decisions that have already been made than for decisions that are yet to be made. Moreover, after making a decision, people tend to recall the positive aspects of the chosen option and the negative aspects of the forgone option (Mather, Shafir, and Johnson, 2003).

[7] Barberis et al (1998) argue that the conservatism bias may explain the underreaction of stock prices to news.

[8] See Kunda (1990) for a summary of the literature.

Relatedly, several studies document an "asymmetric attribution after success and failure"(Gollwitzer, Earle, and Stephan, 1982, Zuckerman, 1979). Shepperd et al. (2008) summarize the findings as follows:

> Several decades of research document a consistent asymmetry in the attributions people make for their personal outcomes. In general, people make internal attributions for desired outcomes and external attributions for undesired outcomes. (...) [This] bias occurs for a variety of events and in a variety of settings. It is evident in workers who attribute receiving promotions to hard work and exceptional skill, yet attribute denial of promotions to unfair bosses. It is evident among athletes who are more likely to assume personal responsibility when they perform well in the sports arena than when they perform poorly (...). It is even evident in drivers who attribute accidents to external factors - the weather, the condition of their car, other drivers - yet attribute the narrow avoidance of an accident to their alertness and finely honed driving skills.

Consistently with these findings, Propositions 3 and 4 show that preferences affect beliefs in the selective awareness model: Individuals with the same prior beliefs and subject to the same information may hold systematically different posterior beliefs if they have different payoffs associated with the information.

3 The Model

The model considers an individual who observes a sequence of independent and identically distributed signals before making a decision. Each signal $\sigma_t \in \{H, L\}$ provides information about the individual's unknown skills $\theta \in [0, 1]$. For simplicity, I assume that the individual has a uniform prior distribution over the unit interval for θ (Subsection 4.1 considers general prior distributions).[9]

After each period, there is a small probability $\eta > 0$ that the process will terminate and the individual will have to take an action $a \in A$, where A is a non-empty compact subset of the Euclidean space.[10] The payoff from action $a \in A$ is determined by $V(a, \theta)$, where $V(., \theta) : A \to \mathbb{R}$ is a continuous function.

The individual processes information as in the models of Benabou and Tirole (2002, 2004, 2006a, 2006b) and Benabou (2008a, 2008b). After observing each signal, the individual decides how to record it. Thus, she chooses whether to interpret the signal realistically $\hat{\sigma}_t = \sigma_t$ or to rationalize it as a different signal $\hat{\sigma}_t \neq \sigma_t$. Figure 1 presents the informational structure in each period.

Because the individual remembers her interpretations of the signal but not the signal itself, a period t history is a sequence of interpretations $h^t = \{\hat{\sigma}_1, \hat{\sigma}_2, ..., \hat{\sigma}_{t-1}\}$, where we define the initial history to be the null set, $h^1 = \varnothing$. When the process terminates, the individual has to take an action $a \in A$ conditional on the current history. Let $\mathcal{H}^t \equiv \{L, H\}^{t-1}$ denote the set

[9] Since I have not assumed that the individual holds a 'correct' prior distribution over θ, she is allowed to hold optimistic or pessimistic beliefs.

[10] Section 4.3 considers the case in which the individual observes a fixed number of signals. Appendix A considers a model in which the individual takes an action in each period.

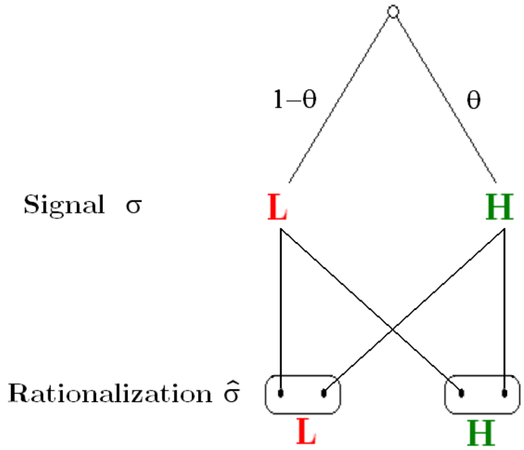

Figure 1: Informational Structure

of period t histories and denote the set of all possible histories by

$$\mathcal{H} \equiv \bigcup_{t=1}^{\infty} \mathcal{H}^t.$$

For any $t' \leq t$, the sequence $h^{t'} = \{\hat{\sigma}_1, \hat{\sigma}_2, ..., \hat{\sigma}_{t'-1}\}$ is said to be a subhistory of $h^t = \{\hat{\sigma}_1, \hat{\sigma}_2, ..., \hat{\sigma}_{t-1}\}$. We write $h^{t'} \subset h^t$ if $h^{t'}$ is a subhistory of h^t.

Following Benabou and Tirole (2006b) and Gottlieb (2009), I assume that the individual has preferences over her skills θ.[11] A high signal $\sigma = H$ is good news about the individual's skills. For simplicity, I consider the following parametrization, which is also generalized in Subsection 4.1:

$$\Pr(\sigma = H|\theta) = \theta.$$

Preferences are additively separable between skills and actions, and skills are normalized to be measured in terms of payoffs. Thus, in each period the individual obtains a payoff from skills equal to the expected value of θ given her interpretations. Payoffs from skills are discounted at rate $\beta \in [0,1]$. Thus, the individual's discounted payoff from skills given history h^t is

$$[1 - \beta(1-\eta)] \sum_{s=0}^{\infty} \beta^s E\left[E\left[\theta|h^{t+s}\right]|h^t\right], \tag{1}$$

where the first expectation is with respect to the probability of reaching history h^{t+s}, which will be a function of the agent's interpretations (endogenous) and the probability of the game ending (exogenous).[12] The term $[1 - \beta(1-\eta)]$ normalizes the sum of discounted payoffs to

[11] Although it is natural to interpret θ as representing the individual's skills, it can be any payoff-relevant characteristic that is positively correlated with the payoff from the action. Subsection 4.2 substitutes this specification by preferences that feature anticipatory utility.

[12] Since $\eta > 0$, the sum in Equation (1) converges even when the individual does not discount the future (i.e., $\beta = 1$).

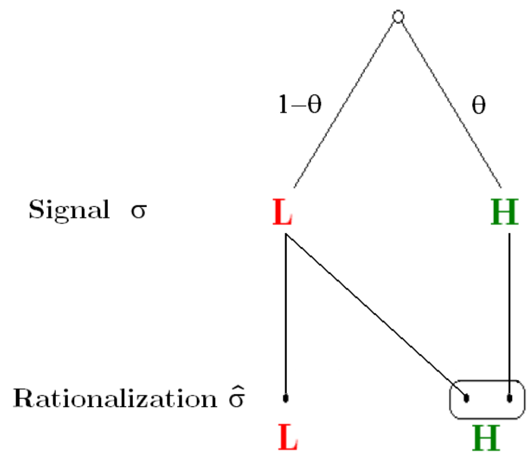

Figure 2: Revised Informational Structure (Benabou-Tirole)

be measured in units of per-period payoffs. Thus, an individual with skill θ has a utility function

$$\alpha\left[1-\beta\left(1-\eta\right)\right]\sum_{s=0}^{\infty}\beta^{s}E\left[E\left[\theta|h^{t+s}\right]|h^{t}\right]+V\left(a,\theta\right),$$

where $\alpha>0$ parametrizes the importance of self-image.[13]

The resulting decision problem features imperfect recall. I follow Piccione and Rubinstein (1997) in modeling a decision problem with imperfect recall as a game between different selves. The individual is treated as a collection of selves, each of them unable to control the behavior of future selves. In each stage before the process terminates, a new self receives a signal σ_t and chooses how to interpret it. When the process terminates, a new self takes an action $a\in A$.

Because we will consider pure strategies only and the signals are ordered by stochastic dominance, we can assume that the individual always interprets high signals realistically (up to a relabeling of interpretations $\hat{\sigma}_t$).[14] Then, after observing a low signal $\sigma_t=L$, she must choose whether to interpret it realistically $\hat{\sigma}_t=L$ or to rationalize it as a high signal $\hat{\sigma}_t=H$ (see Figure 2).

A strategy is a mapping $(\hat{\sigma},a):\mathcal{H}\to\{L,H\}\times A$ determining how to interpret a low signal if the game does not end and which action to take if the game ends. We refer to $\hat{\sigma}:\mathcal{H}\to\{L,H\}$ as an "interpretation strategy" and to $a:\mathcal{H}\to A$ as an "action strategy."

Let $\mu\left(.|h^t\right)$ denote the individual's posterior beliefs about θ given h^t and let $E_\mu\left[.|h^t\right]$ denote the expectation operator with respect to $\mu\left(.|h^t\right)$. Given an action strategy $a:\mathcal{H}\to A$ and posterior beliefs μ, let $\tilde{V}\left(a;h^t,\hat{h}^t\right)$ denote the expected payoff from actions conditional on history h^t when the individual follows the actions prescribed by a in the continuation

[13] I assume that the individual does not discount the payoff from actions because I do not want the payoff from actions to converge to zero as the expected length of the sequence grows. The results in this paper continue to hold if we assume that the individual discounts the payoff from actions at any rate $\delta\in[0,1]$.

[14] See Appendix B for a formal proof of this claim. Appendix A.2 considers mixed strategy equilibria.

histories following $\hat{h}^t$:

$$\begin{aligned}\tilde{V}\left(a;h^t,\hat{h}^t\right) &\equiv E_\mu\left[V\left(a\left(\hat{h}\right),\theta\right)|h^t\right] \\ &= \eta\sum_{s=0}^{\infty}(1-\eta)^s E_\mu\left[V\left(a\left(\hat{h}^{t+s}\right),\theta\right)|h^t\right].\end{aligned}$$

The individual's decision is modeled as the perfect Bayesian equilibrium (PBE) of the multiself game.

Definition 1 *A PBE of the game is a strategy profile* $(\hat{\sigma}^*, a^*): \mathcal{H} \to \{L,H\} \times A$ *and posterior beliefs* μ *such that, for all* $h^t \in \mathcal{H}$,

1. $\hat{\sigma}^*\left(h^t\right) \in \arg\max_{\hat{\sigma}\in\{L,H\}}\left\{\alpha\left[1-\beta\left(1-\eta\right)\right]\sum_{s=0}^{\infty}\beta^s E\left[E_\mu\left[\theta|h^{t+s}\right]|h^t,\hat{\sigma}\right]+\tilde{V}\left(a^*;\left(h^t,L\right),\left(h^t,\hat{\sigma}\right)\right)\right\}$
2. $a^*\left(h^t\right) \in \arg\max_{a\in A}\left\{E_\mu\left[\alpha\theta+V\left(a,\theta\right)|h^t\right]\right\}$
3. $\mu\left(.|h^t\right)$ *is obtained by Bayes' rule if* $\Pr\left(h^t|\hat{\sigma}^*\right)>0$

Conditions 1 and 2 are the perfection conditions. Condition 1 states that each self chooses the interpretation strategy that maximizes her expected payoff given the interpretation and action strategies of other selves. Condition 2 states that, for each terminal history, the individual takes the action that maximizes her expected payoff. Condition 3 is the consistency condition, requiring beliefs to satisfy Bayes' rule given the equilibrium strategies.[15] Given a PBE $(\hat{\sigma}^*, a^*, \mu)$, we say that history h^t is *on the equilibrium path* if $\Pr\left(h^t|\hat{\sigma}^*\right)>0$.

Suppose the individual rationalizes every low signal away, assigns the same posterior as her prior distribution to both interpretations in every period, and chooses a preferred action given her prior beliefs. Because posteriors are not affected by her interpretations, both the interpretation strategy and the action strategy are optimal given beliefs. Moreover, since interpretations are uninformative, beliefs are consistent with Bayes' rule (on the equilibrium path). Hence, there always exists a PBE in which the individual rationalizes away every low signal:

Proposition 1 *There exists a PBE in which* $\hat{\sigma}_t^*\left(h^t\right)=H$ *for all* $h^t\in\mathcal{H}$.

Given a strategy $(\hat{\sigma}^*, a^*)$ and posterior beliefs μ consistent with this strategy, we say that a history h^t is *informative* if the individual's posterior beliefs are affected by its last component:

$$\mu\left(\theta|h^t\right)\neq\mu\left(\theta|h^{t-1}\right).$$

Similarly, a subhistory h^τ of h^t is informative if $\mu\left(\theta|h^\tau\right)\neq\mu\left(\theta|h^{\tau-1}\right)$.

[15] Note that this definition of PBE does not require the individual to assign probability 1 to signal $\sigma = L$ upon observing $\hat{\sigma} = L$ when the equilibrium strategy assigns $\hat{\sigma}^* = H$ because the consistency requirement only applies to actions on the equilibrium path. Therefore, this is an extremely weak definition of a PBE, as it does not even require subgame perfection for games of complete information. When analyzing games of incomplete information, one typically imposes additional restrictions on posterior beliefs (see, e.g., Fudenberg and Tirole, 1991 pp. 331-333). Nevertheless, since the results presented here hold for all PBE, it suffices to work with this less restrictive definition.

Bayesian updating implies that a recollection affects the individual's beliefs only when the individual would have interpreted a low signal realistically, i.e., $\hat{\sigma}^*\left(h^t\right)=L$. If the individual interprets both signals as a high signal $\left(\hat{\sigma}^*\left(h^t\right)=H\right)$, then observing a high recollection is not informative. Thus, when beliefs are consistent with the interpretation strategy, histories in which $\hat{\sigma}^*\left(h^t\right)=H$ do not affect the individual's beliefs: A history h^t is informative if and only if $\hat{\sigma}^*\left(h^t\right)=L$. The following lemma states this result formally:

Lemma 1 *Fix a strategy $(\hat{\sigma}^*, a^*): \mathcal{H} \rightarrow \{L, H\} \times A$ and let μ be posterior beliefs consistent with this strategy. Let $h^t \in \mathcal{H}$ be a history such that $\Pr\left(h^t|\hat{\sigma}^*\right)>0$. h^t is informative if and only if $\hat{\sigma}^*\left(h^t\right)=L$.*

Given a strategy and posterior beliefs consistent with this strategy, we say that a history h^t has k high interpretations in n informative subhistories if k out of the n informative subhistories feature a high interpretation as the last component.[16] The following examples illustrate this definition:

Example 1 *For all $h^t \in \mathcal{H}$, let $\hat{\sigma}^*\left(h^t\right)=H$ and let posterior beliefs $\mu\left(.|h^t\right)$ be equal to the prior distribution. Only high interpretations occur with positive probability. Furthermore, any history $h^t=\{H, ..., H\}$ has 0 high interpretations in 0 informative subhistories.*

Example 2 *Let $\hat{\sigma}^*\left(h^t\right)=L$ for all $h^t \in \mathcal{H}$, and let μ be beliefs consistent with this interpretation strategy. Since every history is informative, a history h^t features $k=\#\left\{\hat{\sigma} \in h^t : \hat{\sigma}=H\right\}$ high interpretations in $n=t$ informative subhistories.*

Example 3 *Let $\hat{\sigma}^*(\varnothing)=L$ and $\hat{\sigma}^*\left(h^t\right)=H$ for all $h^t \neq \varnothing$. Let μ be beliefs consistent with this interpretation strategy. Then, history $\{H, L, L, L..., L\}$ has 1 high interpretation in 1 informative subhistory, and history $\{L, L, L, ..., L\}$ has 0 high interpretations in 1 informative subhistory.*

The conservatism bias states that individuals update beliefs in the right direction, but by too little relative to the Bayesian update. The individual in this model always displays a conservatism bias. When the equilibrium assigns a low interpretation to a low signal $\hat{\sigma}^*\left(h^t\right)=L$, the interpretation fully reveals which signal was observed and beliefs about θ are updated according to Bayes' rule. However, when the equilibrium assigns a high interpretation to a low signal $\hat{\sigma}^*\left(h^t\right)=H$, the individual's recollection is uninformative and, therefore, she does not update her beliefs. Hence, selective interpretation introduces additional noise in the individual's recollections of signals, which leads her to update beliefs in the same direction as the Bayesian update conditional on the realized signals, but at a slower rate.

Let θ_t^B denote the expected value of θ obtained by Bayes' rule conditional on the sequence of observed signals $\{\sigma_1, ..., \sigma_{t-1}\}$, and let $\hat{\theta}_t$ denote the expected value of θ calculated according to the individual's beliefs μ (i.e., conditional on the sequence of the individual's

[16] More precisely, given a strategy and posterior beliefs consistent with this strategy, a history h^t has k high interpretations in n informative subhistories if

$$n=\#\left\{h^{t'} \subset h^t : \mu\left(.|h^\tau\right) \neq \mu\left(.|h^{\tau-1}\right)\right\}, \text{ and } k=\#\left\{\hat{\sigma}_{t'-1}=H : \left(\hat{\sigma}_1, ..., \hat{\sigma}_{t'-1}\right) \subset h^t, \mu\left(.|h^\tau\right) \neq \mu\left(.|h^{\tau-1}\right)\right\}.$$

interpretations). The following proposition establishes that the individual displays conservatism: Posterior expectations move in the same direction, but are "less variable" (in the sense of second-order stochastic dominance) than the expectations obtained by Bayes' rule.

Proposition 2 (Conservatism) *Let $\hat{\sigma}^*$ be an interpretation strategy and let μ be posterior beliefs consistent with this strategy. For any history h^t such that $\Pr\left(h^t|\hat{\sigma}^*\right) > 0$,*

$$\begin{aligned} \theta_t^B &> \theta_{t-1}^B \implies \hat{\theta}_t \geq \hat{\theta}_{t-1}, \text{ and} \\ \theta_t^B &< \theta_{t-1}^B \implies \hat{\theta}_t \leq \hat{\theta}_{t-1}. \end{aligned}$$

Furthermore, $\hat{\theta}_t$ second-order stochastically dominates θ_t^B.

3.1 Binary Actions

This subsection considers a version of the model with a binary action space and symmetric payoffs. The payoff from action $a \in \{0,1\}$ is

$$V(a,\theta) = \begin{cases} \kappa a \text{ if } \theta \geq \frac{1}{2} \\ -\kappa a \text{ if } \theta < \frac{1}{2} \end{cases}, \tag{2}$$

where the parameter $\kappa > 0$ captures the importance of taking the correct action (see Subsubsection 4.1.2 for more general payoff functions). The action can be interpreted as a decision to invest in a project. The project leads to a profit of κ if the individual is sufficiently skilled $\left(\theta \geq \frac{1}{2}\right)$. Otherwise, it leads to a loss of κ.

The following lemma establishes the action taken when the game ends. The symmetry of the payoff function combined with the uniform prior leads to a simple optimal action strategy in which the high action $a = 1$ is taken if the proportion of high recollections in informative subhistories is greater than $\frac{1}{2}$:

Lemma 2 *Fix a PBE $(\hat{\sigma}^*, a^*, \mu)$. Let h^t be a history such that $\Pr\left(h^t|\hat{\sigma}^*\right) > 0$ and suppose that h^t has k high recollections in n informative subhistories. Then,*

$$a^*\left(h^t\right) = \begin{cases} 1 \textit{ if } \frac{k}{n} > \frac{1}{2} \\ 0 \textit{ if } \frac{k}{n} < \frac{1}{2} \end{cases}.$$

Next, we study the properties of Perfect Bayesian Equilibria (3.1.1) and Markovian Equilibria (3.1.2).

3.1.1 Perfect Bayesian Equilibria

When deciding whether to interpret a low signal as a high signal, the individual balances the gain from having a higher self-image with the cost from possibly taking a worse action. The cost of taking a worse action is proportional to the probability that the rationalized signal becomes pivotal in the decision. After either a sufficiently positive or a sufficiently negative sequence of signals, the probability of each signal affecting the decision becomes arbitrarily small. Therefore, the individual will ignore every additional information after such a sequence. This result is formally established in the following theorem:

Theorem 1 (Weak Confirmation Bias with Binary Actions) *Consider the model with binary actions. There exists* $\bar{\eta} > 0$ *such that whenever* $\eta \in (0, \bar{\eta})$:

1. *There exist* $\gamma_L(\eta) > 0$ *and* $\bar{n}(\eta) \in \mathbb{N}$ *such that, in any PBE, no histories with k high interpretations in n informative sub-histories are informative (i.e.,* $\hat{\sigma}^*\left(h^t\right) = H$*) if* $n \geq \bar{n}(\eta)$ *and* $\frac{k}{n} < \gamma_L(\eta)$; *and*

2. *There exist* $\gamma_H(\eta) < 1$ *and* $\bar{n}(\eta) \in \mathbb{N}$ *such that, in any PBE, no histories with k high interpretations in n informative sub-histories are informative (i.e.,* $\hat{\sigma}^*\left(h^t\right) = H$*) if* $n \geq \bar{n}(\eta)$ *and* $\frac{k}{n} > \gamma_H(\eta)$.

Theorem 1 states that the individual disregards all new information after both a sufficiently positive and a sufficiently negative sequence of signals. Therefore, the desire to have higher self-views leads the individual to ignore information after she is sufficiently convinced of which action to take. Whereas this effect has been discussed in the psychology literature for positive information, it may seem somewhat paradoxical that a desire to have higher self-views would lead the individual to be caught in a pessimistic trap, in which attempts to disregard negative information are fruitless since the individual ends up doubting her recollections.

Nevertheless, this result is consistent with a growing literature in psychology relating self-esteem and attributional biases. According to this literature, individuals sufficiently low in self-esteem or suffering from depression display a different attributional style from other individuals. Whereas most individuals tend to attribute success to internal factors (e.g., their own ability and effort), and tend to attribute failure to external factors (e.g., bad luck), low self-esteem or depressed individuals tend to display the opposite pattern.[17] When presented with identical information, low self-esteem and depressed individuals tend be overly pessimistic whereas other individuals tend to hold overly optimistic beliefs.[18] Indeed, it has been argued that the presence of individuals with the opposite pattern of attributions poses an important challenge for theories of motivated cognition (e.g., Moore, 2007). Theorem 1 illustrates that this pattern emerges as a side effect of motivated cognition when individuals are aware of their attempts to improve their self-views.

3.1.2 Markovian Equilibria

Following Maskin and Tirole (2001), I define a Markovian strategy as a strategy that only depends on payoff-relevant information. The only way in which the individual's interpretation of a signal $\hat{\sigma}$ affects her payoff from self-image is through her beliefs about θ. Similarly, the expected payoff from actions depends on the individual's interpretations only through her beliefs about θ. Therefore, a Markovian strategy space partitions the set of histories based on the number of high interpretations and informative subhistories.

More precisely, for any histories h^t and $h^{\tau'}$ with k high interpretations in n informative states, we have $\mu\left(.|h^t\right) = \mu\left(.|h^{\tau'}\right)$. Since the individual's actions affect her expected payoff only through beliefs, it follows that a minimal sufficient partition groups together every

[17] See e.g. Kuiper (1978), Tennen and Herzberger (1987), Peterson et al. (1981), Seligman eta l. (1979), and references therein.

[18] See Beck (1967), Alloy and Ahrens (1987) and references therein.

history with the same number of high interpretations k and informative states n. Given a strategy and beliefs consistent with this strategy, let $\phi : \mathcal{H} \to \left\{(k,n) \in \mathbb{N}^2 : k < n\right\}$ denote the mapping that associates each history $h^t \in \mathcal{H}$ to its number of high interpretations and informative subhistories (k,n). A Markovian Perfect Bayesian equilibrium is a PBE with the property that strategies depend only on the state (k,n) :

Definition 2 *A Markovian Perfect Bayesian equilibrium (MPBE) of the game is a PBE such that for any histories h^t and $h^{\tau'}$ in $\mathcal{H}$, $\phi\left(h^t\right) = \phi\left(h^{\tau'}\right) \implies \hat{\sigma}\left(h^t\right) = \hat{\sigma}\left(h^{\tau'}\right)$ and $a\left(h^t\right) = a\left(h^{\tau'}\right)$.*

Remark 1 *This definition follows Maskin and Tirole in excluding the history length from the state space. Some authors refer to these strategies as stationary Markovian strategies (e.g., Bhaskar, Mailath, and Morris, 2009).*

Remark 2 *Given an MPBE, a Markovian strategy $\hat{\sigma}_M \equiv \hat{\sigma} \circ \phi$ and $a_M \equiv a \circ \phi$ associates each state to the action taken by the individual in that state. Accordingly, we define the individual's Markovian beliefs as the belief associated with each state: $\mu_M(.|k,n) \equiv \int d\mu\left(.|\phi\left(h^t\right) = (k,n)\right)$.*

Lemma 2 established that in any PBE, action strategies must agree in histories such that either $2k > n$ or $2k < n$. Therefore, the only additional restriction imposed by the Markovian assumption is that they must also agree in histories such that $2k = n$ (i.e., when the individual is indifferent between both actions). The Markovian restriction, however, greatly simplifies the analysis of interpretation strategies $\hat{\sigma}$.

Let $\hat{\sigma}_M^*$ be a Markovian strategy in which the individual interprets a signal realistically at state (k,n), i.e. $\hat{\sigma}_M^*(k,n) = L$. If the game does not end at (k,n), either the individual observes a high signal and moves to state $(k+1, n+1)$ or she observes a low signal and moves to state $(k, n+1)$. Hence, (k,n) is a transient state whenever $\hat{\sigma}_M^*(k,n) = L$.

Now suppose $\hat{\sigma}_M^*$ is such that the individual rationalizes low signals away at state (k,n), i.e. $\hat{\sigma}_M^*(k,n) = H$. Because both signals are interpreted equally, the interpretation of a high signal is not informative and the individual remains in state (k,n) regardless of which signal is observed. Thus, (k,n) is an absorbing state if $\hat{\sigma}_M^*(k,n) = H$. Therefore, in an MPBE, if the individual rationalizes one signal away, she will keep rationalizing every new signal away in the future.

From Proposition 1, there always exists an MPBE in which the individual discards every information on the equilibrium path: $\hat{\sigma}_M^*(0,0) = H$. The following proposition shows that this is the (essentially) unique MPBE when the payoff from the action is small relative to the payoff from self image:[19]

Proposition 3 (Uniqueness) *Suppose $\kappa < \frac{\alpha}{3}\left[1 - \beta(1-\eta)\right]$. Then, in any MPBE the individual discards every information on the equilibrium path, $\hat{\sigma}_M^*(0,0) = H$.*

As the next proposition shows, other MPBEs exist when the payoff from the action is not small relative to the payoff from self image:

[19] The MPBE is essentially unique in the sense that all different MPBEs agree on actions and beliefs along the equilibrium path.

Proposition 4 (Multiplicity) *Suppose $\kappa \geq \frac{2\alpha}{3}$. There exists an MPBE in which the individual discards all information after the first signal: $\hat{\sigma}_M^* (0,0) = L$ and $\hat{\sigma}_M^* (k,n) = H$ for $(k,n) \neq (0,0)$.*

Remark 3 (Preferences affect Beliefs) *Propositions 3 and 4 show that individuals with the same prior beliefs and subject to the same information may hold systematically different posterior beliefs if they have different payoffs from self-image or from actions. This result is consistent with the evidence that preferences may affect beliefs described in Section 2.*

The following theorem extends the confirmation bias result from Theorem 1 to almost every history by focussing on Markovian equilibria. It establishes that the individual will (almost always) ignore additional information after sufficiently long sequences (not only sufficiently good or bad sequences as in Theorem 1).

Theorem 2 (Strong Confirmation Bias with Binary Actions) *Consider the model with binary actions. There exists $\bar{\eta} > 0$ such that whenever $\eta \in (0, \bar{\eta})$, there exists a date $\bar{t}(\theta, \eta) \in \mathbb{N}$ such that no histories h^t on the equilibrium path are informative for any $t > \bar{t}(\theta, \eta)$ in any MPBE for all $\theta \neq \frac{1}{2}$.*

As a direct consequence of Theorem 2, it follows that the individual's beliefs (almost) never converge to the truth as the number of signals grows (i.e., $\eta \to 0$). Therefore, selective awareness imposes a limit to learning. Moreover, any history on the equilibrium path can be split in two stages. In the first stage, signals are interpreted correctly and beliefs evolve according to Bayes' rule conditional on the observed signals. In the second stage, signals are misinterpreted and beliefs remain unchanged. Consequently, the individual attaches a disproportionately high weight to initial information (i.e., first impressions matter). The following corollary sates this result formally:

Corollary 1 (First Impressions Matter) *Consider the model with binary actions. For any MPBE there exists $\bar{\eta} > 0$ such that whenever $\eta \in (0, \bar{\eta})$, for any history h^t on the equilibrium path there exists some $\bar{\tau} \leq t$ such that a subhistory $h^\tau \subset h^t$ is informative if and only if $\tau < \bar{\tau}$.*

3.2 Continuum of Actions

This section considers a version of the model with continuous action spaces. Recall that the action space A is a non-empty compact subset of the Euclidean space. I will make the following assumptions about the payoff from actions $V(a, \theta)$:

Assumption 1. $V(., \theta) : A \to \mathbb{R}$ is continuously differentiable and strictly concave.

Assumption 2. $\arg\max_{a \in A} \{V(a, \theta)\} \in int(A)$ for almost all θ.

As in the binary actions case, the individual balances the gains from self-image with the cost from taking worse actions when deciding how to interpret low signals. Since $E[\theta|k+1, n] - E[\theta|k, n] = \frac{\alpha}{n+2}$, the self-image gain from rationalizing a low signal away is of order $\frac{1}{n}$. However, the cost of distorting actions close to the optimum is of order $\frac{1}{n^2}$. Therefore, for large n, the individual prefers to rationalize every low signal away in any PBE. The next theorem establishes this result formally.

Theorem 3 (Strong Confirmation Bias for Continuous Actions) *Suppose Assumptions 1 and 2 are satisfied. There exists $\bar{\eta}$ such that whenever $\eta \in (0, \bar{\eta})$, there exists an $n_\eta \in \mathbb{N}$ such that, in any PBE, every history on the equilibrium path has at most n_η informative sub-histories.*

Note that, unlike Theorem 1, the statement holds for all histories on the equilibrium path instead of only extreme histories (i.e. when $\frac{k}{n}$ is either sufficiently high or sufficiently low). For Markovian equilibria, Theorem 3 implies that histories on the equilibrium path can be split in two stages. In the first stage, signals are interpreted correctly and beliefs evolve according to Bayes' rule conditional on the observed signals; in the second stage, signals are misinterpreted and beliefs remain unchanged.

Corollary 2 (First Impressions Matter) *Suppose Assumptions 1 and 2 are satisfied. For any MPBE there exists $\bar{\eta} > 0$ such that whenever $\eta \in (0, \bar{\eta})$, for any history h^t on the equilibrium path there exists some $\bar{\tau} \leq t$ such that a subhistory $h^\tau \subset h^t$ is informative if and only if $\tau < \bar{\tau}$.*

4 Generalizations and Extensions

4.1 General Information Structures

In the previous section, I assumed that the individual's prior distribution over skills was uniform and that the payoff from skills was proportional to the expected probability of observing a high signal (i.e., $\Pr(\sigma = H|\theta) = \theta$). In this section, I show that the main results generalize for any regular prior distribution and any conditional distribution of high signals strictly increasing in the individual's skills. Therefore, the main results from Section 3 do not rely on the specific distributional assumptions.

As in Section 3, the model features an individual with preferences over her skills θ. With no loss of generality, we can assume that skills are measured in payoff units. Let the space of possible skills $\Theta = \left(\underline{\theta}, \overline{\theta}\right)$ be a non-empty open interval of the real line and let $\rho(.)$ denote the agent's prior distribution of θ. Denote the probability of observing a high signal conditional on θ by

$$\Pr(\sigma = H|\theta) = \pi(\theta).$$

I assume that $\pi : \left(\underline{\theta}, \overline{\theta}\right) \to [0, 1]$ is a strictly increasing function so that a high signal is more likely under higher skills.[20] Moreover, I also assume that π is twice continuously differentiable and satisfies $\pi(\underline{\theta}) > 0$ and $\pi\left(\overline{\theta}\right) < 1$.

The prior distribution satisfies the following regularity conditions:

Definition 3 *A prior distribution $\rho(.)$ is regular if (i) ρ is thrice continuously differentiable, and (ii) $\rho(\theta) > 0$ for all $\theta \in \Theta$.*

Note that virtually all bounded continuous distributions used in applications satisfy these conditions. The following lemma determines that asymptotic behavior of the conditional expectation $E[\theta|k, n]$ as the number of informative subhistories n increases.

[20] From a statistical perspective, the assumption that π is strictly increasing provides an identification condition for θ. If π were not a one-to-one function, it would be impossible for an individual to learn the true parameter θ regardless of the number of observations.

Lemma 3 *Suppose $\rho(.)$ is regular. There exists constants C_0 and N_k such that*

$$\left| E[\theta|k,n] - \pi^{-1}\left(\frac{k}{n}\right) - \xi\left(\frac{k}{n}\right)\frac{1}{n}\right| \leq \frac{C_0}{n^2} \tag{3}$$

for all $n > N_k$ on an almost sure set (under the true θ), where

$$\xi(x) \equiv \frac{\sqrt{x(1-x)}}{\pi'(\pi^{-1}(x))} \times \left\{ \frac{1}{x(1-x)} \left\{ \begin{array}{c} 2\left[\frac{1-2x}{x(1-x)}\right]\left[\pi'\left(\pi^{-1}(x)\right)\right]^3 \\ -3\pi'\left(\pi^{-1}(x)\right)\pi''\left(\pi^{-1}(x)\right) \end{array} \right\} + \frac{\rho'\left(\pi^{-1}(x)\right)}{\rho\left(\pi^{-1}(x)\right)} \right\} \tag{4}$$

Therefore, for any regular prior distribution ρ, the expected payoff converges to its maximum likelihood estimator $\pi^{-1}\left(\frac{k}{n}\right)$ plus terms of order higher than $\frac{1}{n}$. The next lemma determines the self-image gain from deviating from a history with k high interpretations to a history with $k+1$ high interpretations when the number of informative subshistories n is large.

Lemma 4 *Suppose $\rho(.)$ is regular. There exists constants C_1 and N_k such that*

$$\left| E[\theta|k+1,n] - E[\theta|k,n] - \frac{1}{\pi'\left(\pi^{-1}\left(\frac{k}{n}\right)\right)}\frac{1}{n}\right| \leq \frac{C_1}{n^2},$$

for all $n > N_k$ on an almost sure set (under the true θ).

As in the case of uniform priors with linear conditional probabilities studied in Section 3, the self-image gain of deviating from a history with k high interpretations to a history with $k+1$ high interpretations has order $\frac{1}{n}$. Next, we will study the model with continuous actions and binary actions separately.

4.1.1 Continuous Actions

Consider an equilibrium in which the individual interprets a low signal realistically after history h^t, i.e. $\hat{\sigma}^*\left(h^t\right) = L$. From Lemma 4, the self-image gain from deviating to $\hat{\sigma} = H$ has order $\frac{1}{n}$. However, as in Subsection 3.2, it can be shown that the cost of this deviation in terms of worsened actions is of order $o\left(\frac{1}{n}\right)$. Therefore, when n is large enough, the gain from self-image dominates the loss from worse actions and the individual prefers to rationalize low signals away in any PBE:

Theorem 4 (Strong Confirmation Bias for Continuous Actions) *Suppose $\rho(.)$ is regular and suppose Assumptions 1 and 2 are satisfied. There exists $\bar{\eta}$ such that whenever $\eta \in (0,\bar{\eta})$, there exists an $n_\eta \in \mathbb{N}$ such that, in any PBE, every history on the equilibrium path has at most n_η informative sub-histories.*

The following corollary extends Corollary 2:

Corollary 3 (Strong Confirmation Bias for Continuous Actions) *Suppose $\rho(.)$ is regular and suppose Assumptions 1 and 2 are satisfied. For any MPBE there exists $\bar{\eta} > 0$ such that whenever $\eta \in (0,\bar{\eta})$, for any history h^t on the equilibrium path there exists some $\bar{\tau} \leq t$ such that a subhistory $h^\tau \subset h^t$ is informative if and only if $\tau < \bar{\tau}$.*

4.1.2 Binary Actions

Next, we consider a generalized version of the binary actions model of Subsection 3.1. Let the action space be $A = \{0, 1\}$. The following assumption generalizes the payoff function from equation (2):

Assumption 3. $V(a, \theta) = av(\theta)$, $a \in \{0, 1\}$, where $v : \Theta \to \mathbb{R}$ satisfies

$$E[v(\theta) | \sigma_1, ..., \sigma_t] > 0 \iff E[\pi(\theta) | \sigma_1, ..., \sigma_t] > \lambda$$

for some $\lambda \in \mathbb{R}$, $\forall t, \sigma_t$.

Under the payoff function above, the high action $a = 1$ is optimal if $E[v(\theta) | \sigma_1, ..., \sigma_t] > 0$. When v is increasing, there always exists a threshold $\lambda(\sigma_1, ..., \sigma_t)$ such that $a = 1$ is optimal if and only if the expected probability of observing a high signal is above this threshold:

$$E[v(\theta) | \sigma_1, ..., \sigma_t] > 0 \iff E[\pi(\theta) | \sigma_1, ..., \sigma_t] > \lambda(\sigma_1, ..., \sigma_t).$$

Assumption 3 states that the threshold $\lambda(\sigma_1, ..., \sigma_t)$ is not a function of sequence of signals $\{\sigma_1, ..., \sigma_t\}$. It simplifies the analysis because it implies that the threshold which determines whether a recollection is pivotal does not depend on the history of observations.

Note that the framework from Subsection 3.1 satisfies Assumption 3 with $\pi(\theta) = \theta$, $\lambda = \frac{1}{2}$, and

$$v(\theta) = \begin{cases} \kappa \text{ if } \theta \geq \frac{1}{2} \\ -\kappa \text{ if } \theta < \frac{1}{2} \end{cases}.$$

In general, Assumption 3 places restrictions on the distribution of signals and the payoff function v. When the payoff from actions is an affine function $v(\theta) = b\pi(\theta) + c$, however, this assumption is satisfied regardless of the distribution of signals since

$$E[v(\theta) | \sigma_1, ..., \sigma_t] > 0 \iff E[\pi(\theta) | \sigma_1, ..., \sigma_t] > -\frac{c}{b}.$$

The following theorem generalizes the result from Theorem 2. As in the uniform case, the probability that a signal becomes pivotal converges to 0 faster than $\frac{1}{n}$. Therefore, after a sufficiently long history, the individual will always rationalize low signals away.

Theorem 5 (Strong Confirmation Bias with Binary Actions) *Suppose $\rho(.)$ is regular and suppose Assumption 3 is satisfied. There exists $\bar{\eta} > 0$ such that whenever $\eta \in (0, \bar{\eta})$, there exists a date $\bar{t}(\theta, \eta) \in \mathbb{N}$ such that no histories h^t are informative for any $t > \bar{t}(\theta, \eta)$ in any MPBE for almost all θ.*

Corollary 4 (First Impressions Matter) *Suppose $\rho(.)$ is regular and suppose Assumption 3 is satisfied. For any MPBE there exists $\bar{\eta} > 0$ such that whenever $\eta \in (0, \bar{\eta})$, for any history h^t on the equilibrium path there exists some $\bar{\tau} \leq t$ such that a subhistory $h^\tau \subset h^t$ is informative if and only if $\tau < \bar{\tau}$.*

4.2 Anticipatory Utility

In the previous sections, I have followed Benabou and Tirole (2006b) and Gottlieb (2009) in interpreting the parameter θ as the individual's skills. Signals were assumed to be informative about the individual's skills and therefore provided information about the appropriate action to be taken. This subsection follows Benabou (2008a, 2008b) in considering decision makers with anticipatory utility.

Consider an individual who has to take an action $a \in \{-1, 1\}$. The payoff from each action $V(a, \theta)$ is a function of the (unknown) state of the world θ. As in Section 3, let $\Pr(\sigma = H|\theta) = \theta$, and assume that θ is uniformly distributed in the interval $[0, 1]$. The payoff from actions is determined by

$$V(a, \theta) = a\left(\theta - \frac{1}{2}\right).$$

For example, $a = 1$ can be interpreted as buying an asset and $a = -1$ as selling it. The prior distribution over the payoff from each action is uniform over the interval $\left[-\frac{1}{2}, \frac{1}{2}\right]$. Because V is a linear function of a, the results remain unchanged if we allow for additional actions in the interval $(-1, 1)$. In that case, an action can be interpreted as the proportion of the agent's wealth allocated to the buying or selling the asset.

Let μ denote the individual's posterior beliefs about the state of the world θ. When deciding how to interpret a low signal after a history h^t, the current self takes two terms into account. First, she considers the expected future payoff from actions $E_\mu\left[V\left(a\left(h^{\tilde{t}}\right), y\right)|h^t, L\right]$, where $h^{\tilde{t}}$ denotes the (random) history in which the game ends. Second, she takes into account the *anticipatory utility* $E_\mu\left[V\left(a\left(h^{\tilde{t}}\right), y\right)|h^t, \hat{\sigma}\right]$, where $\hat{\sigma}$ is her interpretation of the signal.

The individual chooses the interpretation $\hat{\sigma} \in \{L, H\}$ that maximizes:

$$[1 - \beta(1-\eta)] \sum_{s=1}^{\infty} \beta^t \left\{E\left[E_\mu\left[V\left(a\left(h^{\tilde{t}}\right), y\right)|h^{t+s}\right]|h^t, \hat{\sigma}\right]\right\} + E_\mu\left[V\left(a\left(h^{\tilde{t}}\right), y\right)|h^t, L\right], \quad (5)$$

where $\beta \in [0, 1]$ captures the relative importance of future anticipatory utility to the current self. When $\beta = 0$, the individual cares only about current anticipatory utility. The definitions of equilibria are analogous to the ones from Definitions 1 and 2, with the substitution of the utility function by (5).

In the model of Section 3, each self balanced the gains from higher self-views with the expected costs of making worse decisions when choosing how to interpret each signal. Then, when the individual was sufficiently confident of which action to take, the self-views effect dominated and she always chose to rationalize low signals away. The anticipatory utility features a similar trade-off. Here, the individual balances the gain in anticipatory utility from believing in a better state of the world with the expected cost of making worse decisions. As in the self-views model, when the individual is sufficiently confident of which action to take, the anticipatory utility effect dominates and low signals are rationalized away. The following theorem establishes this result formally:

Theorem 6 (Weak Confirmation Bias for Anticipatory Utility) *Consider the anticipatory utility model with binary actions. There exists $\bar{\eta}$ such that whenever $\eta \in (0, \bar{\eta})$:*

1. *There exist $\gamma_L(\eta) > 0$ and $\bar{n}(\eta) \in \mathbb{N}$ such that, in any PBE, no histories with k high interpretations in n informative sub-histories are informative if $n \geq \bar{n}(\eta)$ and $\frac{k}{n} < \gamma_L(\eta)$; and*
2. *There exist $\gamma_H(\eta) < 1$ and $\bar{n}(\eta) \in \mathbb{N}$ such that, in any PBE, no histories with k high interpretations in n informative sub-histories are informative if $n \geq \bar{n}(\eta)$ and $\frac{k}{n} > \gamma_H(\eta)$.*

4.3 Known Terminal Period

In the model from Section 3, the information acquisition process terminated in each period with a constant probability η. This setup represents a situation in which the individual expects to receive a long sequence of information, but is unsure about when the decision will have to be made. The current subsection considers a setup in which the individual observes a fixed number of signals T before making the decision. This framework depicts situations in which the individual knows when the decision will have to be made and how much information will be acquired before then.

The discounted payoff from skills given history h^t is $\sum_{s=0}^{T-t} \beta^s E\left[E\left[\theta|h^{t+s}\right]|h^t\right]$, where the first expectation is with respect to the probability of reaching history h^{t+s}. Thus, an individual with skill θ has utility function

$$\alpha \sum_{s=0}^{T-t} \beta^s E\left[E\left[\theta|h^{t+s}\right]|h^t\right] + V(a,\theta), \tag{6}$$

where the payoff from actions $a \in \{0,1\}$ is the same as in Subsection 3.1:

$$V(a,\theta) = \begin{cases} \kappa a \text{ if } \theta \geq \frac{1}{2} \\ -\kappa a \text{ if } \theta < \frac{1}{2} \end{cases}.$$

The PBE definition is the same as the one from Definition 1, with the appropriate substitution of the utility function by (6).

Recall that the decision of whether to interpret a low signal realistically balances the benefit from having a higher self-image against the cost of worse decision-making. Therefore, distorting an interpretation is only costly when the interpretation affects which action will be chosen. This observation allows us to place a bound on the number of signals that can be interpreted realistically, even when the payoff from self image is arbitrarily small.

For example, suppose the individual observes $T = 3$ signals, and consider the individual's choice of how to interpret the last signal. Following a history with 0 high interpretations in 2 informative subhistories, the individual will always choose $a = 0$ regardless of the signal observed in the last period and the individual's interpretation strategy in the last period. Similarly, after a history with 2 high interpretations in 2 informative subhistories, she always chooses $a = 1$ in every continuation history for every interpretation strategy following that history. Thus, because in both cases rationalizing a low signal away has no costs, the individual will always choose to rationalize it away. The following proposition applies a backward induction argument to generalize this argument for any finite number of periods:

Proposition 5 (Weak Confirmation Bias for Finitely Many Periods) *Consider the model with a finite number T of signals. Fix a PBE $(\hat{\sigma}^*, a^*, \mu)$ and let $h^{T-\tau}$ be a history on the* *equilibrium path. Suppose that $h^{T-\tau}$ has k high recollections in n informative subhistories, where $0 \leq \tau \leq T$. If either $\frac{k}{n+\tau} \geq \frac{1}{2}$ or $\frac{k+\tau+1}{n+\tau} < \frac{1}{2}$, then $h^{T-\tau}$ is not informative (i.e., $\hat{\sigma}^*\left(h^{T-\tau}\right) = H$).*

Thus, after the individual becomes convinced of which action to take, she disregards every additional information. In particular, we cannot have equilibria in which all signals are interpreted realistically when $T \geq 3$. Moreover, because the possibility of a signal becoming pivotal is increasing in the number of remaining signals τ, information is more likely to be disregarded as the game approaches its end (i.e., for k and n fixed, the bound determining which signals are rationalized away decreases in τ).

Note that, unlike in Theorem 1, the bounds on the proportion of high and low recollections from Proposition 5 is independent of the relative importance self image (parametrized by α). As the $T = 3$ example showed, the cost of deviating in terms of taking a worse action is equal to zero in some histories. On the other hand, in the model with random termination, this cost is never exactly equal to zero since there is always some probability that the individual will eventually observe a large number of future signals, allowing her to possibly take a better action.

Proposition 5 takes the number of signals T as fixed, whereas the analysis from the previous sections considered the individual's behavior as the expected number of signals grew (i.e., $\eta \to 0$). Recall that the main intuition from the section was that as the individual collects more information, the cost of making worse decisions vanishes at rate $\frac{1}{n^2}$ whereas the benefit from having a higher self-image decreases at rate $\frac{1}{n}$. Therefore, the individual prefers to rationalize low signals away after a sufficiently large number of observations. Since this result does not rely on the randomness of terminal date, it is straightforward to generalize all the results from the previous section to the model with a known terminal date, replacing $\eta \approx 0$ by T large enough. The only results that have to be slightly modified are Propositions 3 and 4, which are established below:

Proposition 6 (Uniqueness) *Suppose $\kappa < \alpha \left[\frac{1}{3} + \frac{\beta}{T+3}\left(\frac{1-\beta^T}{1-\beta}\right)\right]$ if $\beta < 1$ or $\kappa < \alpha \sum_{t=3}^{T} \frac{1}{t}$ if $\beta = 1$. In any MPBE the individual discards every information on the equilibrium path, $\hat{\sigma}^*_M(0,0) = H$.*

Proposition 7 (Multiplicity) *Suppose $\kappa \geq \frac{2\alpha}{3}\left(\frac{1-\beta^T}{1-\beta}\right)$. There exists an MPBE in which the individual discards all information after the first signal: $\hat{\sigma}^*_M(0,0) = L$ and $\hat{\sigma}^*_M(k,n) = H$ for $(k,n) \neq (0,0)$.*

5 Implications for Dynamic Information Transmission

Although the model was presented as representing an individual's decision on how to interpret information, it could alternatively be seen as a model of information disclosure between different individuals. Consider a situation involving experts who provide advice to an uninformed individual who has to make a decision. The individual receiving advice will eventually take an action $a \in A$, whose payoff $V(a,\theta)$ depends on an unknown state of the world θ.

In each period, an expert observes a signal $\sigma_t \in \{H, L\}$ that provides information about the state of the world. The expert decides how to describe the signal to the uninformed individual. For example, the expert may be a consultant who observes information about the profitability of an investment and, in each period, strategically interprets the information to an uninformed investor. The expert may also be an employee who communicates strategically to her employer. In another application of the model, a doctor observes a sequence of tests and decides how to strategically interpret them to a patient, who will then choose one treatment.

Experts care about the payoff $V(a, \theta)$ but also care about the uninformed individual's perception of the state of the world. For example, the state of the world may be informative about the experts' skills and experts may care about conveying to have high skills (either because their payment may formally depend on their perceived skills or because of career concerns). Alternatively, doctors may favor states of the world associated with more expensive treatments. Similarly, consultants may favor states of the world in which their advice is perceived as being more important.

When $\beta = 0$, the model is formally equivalent to a cheap-talk game in which a sequence of experts with reputation concerns provide advice to the uninformed individual.[21] For $\beta > 0$, the model differs from a more standard cheap-talk game in that the expert can only access the information transmitted to the uninformed individual but not the actual observed information. Nevertheless, it can be shown that the main results from this paper generalize to the case where the expert can condition her actions on the actual information.

The model's implications can then be stated in terms of learning by the individual receiving the advice. For binary actions, Theorem 1 implies that in any PBE the expert's additional advice is always discarded after a sufficiently positive or sufficiently negative sequence of informative signals. Theorem 2 states that the expert's advice can be partitioned in two stages in any Markovian equilibria. In the first stage, she reveals information truthfully and her advice is correctly interpreted. In the second stage, her additional advice is uninformative. Proposition 3 states that if the expert cares enough about signaling a high state of the world θ, any Markovian equilibrium is fully uninformative. Theorem 3 shows that when actions are continuous, every information gets discarded after a certain number of periods in any PBE. Proposition 5 establishes bounds on the amount of informative advice when the number of signals is finite.

6 Conclusion

Several papers have recently used the selective awareness framework proposed by Benabou and Tirole to provide explanations for deviations from rational decision making. However, it is often argued that the Bayesian updating assumption embedded in this framework combined with the repeated nature of the decisions being modeled would lead individuals to eventually learn the truth and, therefore, the departures from rationality would vanish in the long-run.

This paper formally studied this issue by studying a repeated version of the selective awareness model of Benabou and Tirole (2002, 2004) and Benabou (2008a, 2008b). It showed that all information gets disregarded after a certain number of observations. Therefore, learning is always incomplete, and the departures from rationality presented in the static

[21] Morgan and Stocken (2003) present a (one-period) reputational cheap talk model in which an analyst may have an incentive to misreport her information in order to inflate the stock price. Ottaviani and Sørensen (2006b) consider a related strategic communication model in which a sequence of experts are concerned about appearing to be well informed.

models in the literature do not disappear even when the decision problem is repeated infinitely many times.

The model predicts a behavior that is consistent with some biases in information processing studied by psychologists. Individuals attribute a disproportionately large weight to initial information. After becoming sufficiently convinced of which action to take, they do not change their beliefs (confirmation bias). They also update beliefs in the right direction, but in insufficient amount compared to the Bayesian updating rule (conservatism bias).

Appendix A. Extensions

A.1 Repeated Actions

In the model considered in the text, the individual faced a single action at the end of the game. Because the gain from self-image was greater than the loss from distorting actions after a sufficiently large sequence of observations, any equilibrium featured incomplete learning. The intuition carries over to models in which the individual faces a sequence of actions as long as the actions themselves are not fully informative about the individual's skills. This subsection provides a simple extension of the binary and continuous action models from Section 3 that incorporates an action being taken in each period.

The model considers an infinitely lived individual who takes an action $a \in A$ in each period. After the action is chosen, the individual observes a signal $\sigma \in \{L, H\}$, where

$$\Pr(\sigma = H | a, \theta) = \theta \text{ for all } a \in A. \tag{7}$$

Note that a high signal $\sigma = H$ is good news about the individual's skills θ. Moreover, for simplicity, (7) assumes that any action leads to the same distribution of signals.

As in the general framework of Section 3, the individual's unknown skills are uniformly distributed in the unit interval and, for notational simplicity, we assume that the game is infinitely repeated (i.e., there is no random termination as in the model from Section 3). After observing the realization of the signal σ, the individual forms an interpretation $\hat{\sigma} \in \{L, H\}$. As in the basic framework, there is no loss of generality in assuming that high signals are always interpreted correctly, whereas low signals can be reinterpreted as either high or low.

If a signal σ occurs, the individual obtains a payoff from actions equal to $V(\sigma, a)$. Furthermore, the individual also obtains a payoff from self-image equal to $E_\mu[\theta | h^t]$. Payoffs are discounted at rate β. Hence, the individual's discounted sum of expected payoffs is:

$$(1-\beta) \sum_{s=0}^{\infty} \beta^s \left\{ \alpha E\left[E\left[\theta | h^{t+s} \right] | h^t \right] + V(\sigma, a_{t+s}) \right\}.$$

First, let us consider the binary actions case: $A = \{0, 1\}$. The payoff from action $a \in \{0, 1\}$ is

$$V(\sigma, a) = \begin{cases} \kappa a \text{ if } \sigma = H \\ -\kappa a \text{ if } \sigma = L \end{cases}.$$

A high signal leads to a gain of κ if the individual has taken a high action $a = 1$. Symmetrically, a low signal is bad news about the individual's skills and leads to a loss of κ if she has taken a high action $a = 1$.

Fix a PBE $(\hat{\sigma}^*, a^*, \mu)$. Following the same steps as Lemma 2, it follows that for any history on the equilibrium path with k high interpretations in n informative subhistories, we must have

$$a^*(h^t) = \begin{cases} 1 \text{ if } \frac{k}{n} > \frac{1}{2} \\ 0 \text{ if } \frac{k}{n} < \frac{1}{2} \end{cases}.$$

As in the case of a single action taken at the end of the game, the cost of distorting actions in each period converges to zero faster than $\frac{1}{n}$ whereas the gain from self image converges at rate $\frac{1}{n}$. Therefore, following the exact same steps as Theorems 1 and 2, it is straightforward to establish the following results:

Theorem 7 (Repeated Binary Actions) *Consider the model with repeated binary actions.*

1. *There exists a date $\bar{t}(\theta) \in \mathbb{N}$ such that no histories h^t are informative for any $t > \bar{t}(\theta)$ in any MPBE for all $\theta \neq \frac{1}{2}$;*
2. *There exist $\gamma_L > 0$ and $\bar{n} \in \mathbb{N}$ such that, in any PBE, no histories with k high interpretations in n informative sub-histories are informative if $n \geq \bar{n}$ and $\frac{k}{n} < \gamma_L$; and*
3. *There exist $\gamma_H < 1$ and $\bar{n} \in \mathbb{N}$ such that, in any PBE, no histories with k high interpretations in n informative sub-histories are informative (i.e., $\hat{\sigma}^*(h^t) = H$) if $n \geq \bar{n}$ and $\frac{k}{n} > \gamma_H$.*

In the continuous actions case, it is also straightforward to extend the arguments from Theorem 3 and Corollary 2 to the case of repeated actions. Then, we have the following results:

Theorem 8 (Repeated Continuous Actions) *Suppose Assumptions 1 and 2 are satisfied.*

1. *There exists an $n \in \mathbb{N}$ such that, in any PBE, every history on the equilibrium path has at most n informative sub-histories.*
2. *there exists $\bar{t} \in \mathbb{N}$ such that no histories h^t are informative for any $t > \bar{t}$.*

A.2 Mixed Strategies

In the text, we have focused on pure strategy equilibria. This subsection extends the incomplete learning for mixed strategy equilibria. Consider the continuous action case under quadratic payoffs:

$$V(a, \theta) = -\kappa (a - \theta)^2 . \tag{8}$$

As in Subsubsection 4.1.1, we allow for any regular prior distribution ρ but, for simplicity, we consider $\pi(\theta) = \theta$.

Once we allow for mixed strategies, the equivalence between the decision problem from Figure 1 and the Benabou and Tirole model (Figure 2) may no longer hold. In this Subsection, we will focus on the Benabou and Tirole model.

A mixed strategy are mappings $\alpha : \mathcal{H} \to [0,1]$ and $\gamma : \mathcal{H} \to \Delta(\mathbb{R})$, where $\alpha(h^t)$ assigns a probability of playing $\hat{\sigma} = H$ after observing a low signal given history h^t, and γ assigns the probability of playing each action $a \in \mathbb{R}$ if the game ends after history h^t.

We denote the strategy $\gamma(h^t)$ that assigns probability one to action $a \in \mathbb{R}$ by $\gamma(h^t) = \delta_a$. Under quadratic payoffs (equation 8), any PBE features pure strategy actions:

Lemma 5 *Fix a PBE. Then, $\gamma^*(h^t) = \delta_{E[\theta|h^t]}$.*

Proof. Consider the action that maximizes the expected payoff from actions:

$$\max_a \left[-\kappa \int (a - \theta)^2 f\left(\theta | h^t\right) d\theta \right].$$

The unique solution to this program is $a^* = E[\theta|h^t]$. Since the solution is unique, the individual cannot play mixed action strategies in equilibrium. ■

Lemma 6 *Fix a PBE. For any history h^t on the equilibrium path, we have:*

$$E\left[V\left(\gamma^*\left(h^t,L\right),\theta\right)|h^t,L\right]-E\left[V\left(\gamma^*\left(h^t,H\right),\theta\right)|h^t,L\right]=\kappa\left\{E\left[\theta|h^t,H\right]-E\left[\theta|h^t,L\right]\right\}^2. \qquad (9)$$

Proof. For notational simplicity, for a fixed history h^t, let $\theta_H \equiv E\left[\theta|h^t,H\right]$ and $\theta_L \equiv E\left[\theta|h^t,L\right]$. Substituting the $\gamma^*\left(h^t\right)=\delta_{E[\theta|h^t]}$ in the quadratic payoff function, we obtain

$$\begin{aligned}
& E\left[V\left(\gamma^*\left(h^t,L\right),\theta\right)|h^t,L\right]-E\left[V\left(\gamma^*\left(h^t,H\right),\theta\right)|h^t,L\right] \\
= & -\kappa\left[\int\left(\theta_H-\theta\right)^2 f\left(\theta|h^t,L\right)d\theta-\int\left(\theta_L-\theta\right)^2 f\left(\theta|h^t,L\right)d\theta\right] \\
= & -\kappa\left(\theta_H-\theta_L\right)^2.
\end{aligned}$$

■

Therefore, after history h^t, the cost of taking the action associated with $\sigma_t = H$ if the true signal was $\sigma_t = L$ is proportional to the square of the difference between the conditional expectations of θ in each history. Whenever $E\left[\theta|h^t,H\right]-E\left[\theta|h^t,L\right]<\frac{\alpha}{\kappa}$, it follows that the cost of taking a suboptimal action is smaller than the gain from improving self-image:

$$\alpha E\left[\theta|h^t,H\right]+E\left[V\left(\gamma^*\left(h^t,H\right),\theta\right)|h^t,L\right]>\alpha E\left[\theta|h^t,L\right]+E\left[V\left(\gamma^*\left(h^t,L\right),\theta\right)|h^t,L\right].$$

Therefore, whenever the individual's beliefs is sufficiently close to the true parameter, she prefers to interpret every additional signal as $\hat{\sigma}=H$. Consequently, no additional signal affects posteriors after that point. Thus, beliefs cannot converge to the truth:

Proposition 8 *Consider the continuous actions model with quadratic payoffs. In any (mixed strategy) PBE,*

$$\Pr\left(\lim_{t\to\infty}E\left[\theta|h^t\right]=\theta\right)=0.$$

Before presenting the proof, let us introduce the following notation. An *outcome path* for an infinitely repeated game is an infinite sequence of interpretations $h^\infty \equiv \{\hat{\sigma}_1,\hat{\sigma}_2,...\}\in\{L,H\}^\infty$. A period t history induced by the outcome path h^∞ consists of the $t-1$ first elements of h^∞.

Proof. Fix a PBE. Consider an outcome path h^∞ such that, for every $t\in\mathbb{N}$, the history induced by h^∞ is on the equilibrium path. Suppose that $\lim_{t\to\infty}E\left[\theta|h^t\right]=\theta$. Then, h^∞ must have an infinite number of informative subhistories.

Note that for any history h^t induced by the outcome path h^∞, any history that agrees with h^t in all but one interpretation is on the equilibrium path if and only if the disagreement occurs at an informative subhistory. Therefore, there are infinitely many on-the-equilibrium-path histories which agree with h^∞ in all but one interpretation. Thus, h^∞ cannot have a strictly positive probability mass.

Let $\tilde{h}^\infty$ be an outcome path that agrees with h^∞ in all but one interpretation that happens in an informative subhistory (we have shown that there exist infinitely many of those) and suppose that $\lim_{t\to\infty}E\left[\theta|\tilde{h}^t\right]=\theta$. Since the disagreement occurs in an informative subhistory, it follows that histories induced by h^∞ are also on the equilibrium path.

By definition of convergence, for any $\varepsilon>0$ there exists $\bar{t}$ such that $t>\bar{t}$ implies

$$\left|E\left[\theta|h^t\right]-\theta\right|<\frac{\varepsilon}{2} \quad \text{and} \quad \left|E\left[\theta|\tilde{h}^t\right]-\theta\right|<\frac{\varepsilon}{2}.$$

Thus,

$$\left|E\left[\theta|\tilde{h}^t\right]-E\left[\theta|h^t\right]\right|\leq\left|E\left[\theta|h^t\right]-\theta\right|+\left|E\left[\theta|\tilde{h}^t\right]-\theta\right|<\varepsilon.$$

Take $\varepsilon = \frac{\alpha}{\kappa}$. Then, there exists $\bar{t}$ such that for all $t > \bar{t}$

$$E\left[\theta|h^t, H\right] - E\left[\theta|h^t, L\right] < \frac{\alpha}{\kappa}.$$

But, by equation (9), this implies

$$\alpha E\left[\theta|h^t, H\right] + E\left[V\left(a\left(H\right), \theta\right)|h^t, L\right] > \alpha E\left[\theta|h^t, L\right] + E\left[V\left(a\left(L\right), \theta\right)|h^t, L\right].$$

Therefore, if the disagreement occurs at some $t > \bar{t}$ the individual strictly prefers to play $\hat{\sigma} = H$ (and thus cannot either mix or play L), which implies that h^t and $\tilde{h}^t$ cannot be both on the equilibrium path. Hence, outcome paths $\tilde{h}^\infty$ that agree with h^∞ in all but one interpretation can only feature $\lim_{t\to\infty} E\left[\theta|\tilde{h}^t\right] = \theta$ if the disagreement occurs in a period $t < \bar{t}$. Denote the (finite) set of such paths by $\mathcal{T}$.

Let $\mathcal{P}_{h^\infty}$ be the set of outcome paths with the following properties: (i) every history induced by the outcome path is on the equilibrium path, and (ii) the outcome path agrees with h^∞ in all but one interpretation. As noted previously, this is an infinite set. Since the set $\mathcal{T} \subset \mathcal{P}_{h^\infty}$ is finite and no history has a strictly positive probability mass, it follows that $\Pr\left(\lim_{t\to\infty} E\left[\theta|h^t\right] = \theta\right) = 0$. ■

Appendix B. Auxiliary Results

Lemma 7 *Let $\Pr\left(\sigma = H|\theta\right) = \theta \sim U\left[0,1\right]$ and denote by $f\left(\theta|k,n\right)$ the p.d.f. determined by Bayesian updating conditional on a history with k high recollections in n informative subhistories. Then:*

$$\begin{aligned} f\left(\theta|k,n\right) &= \frac{\theta^k\left(1-\theta\right)^{n-k}}{\int_0^1 \theta^k\left(1-\theta\right)^{n-k} d\theta}, \text{ and} \\ E\left[\theta|k,n\right] &= \frac{k+1}{n+2}. \end{aligned}$$

Proof. Let $x_t = \mathbf{1}\left(\sigma_t = H\right)$. Then, $x_t|\theta$ follows a Bernoulli distribution with parameter θ. By independence, the conditional distribution having k successes in n informative subhistories follows a Binomial distribution

$$\Pr\left(\tilde{k} = k, \tilde{n} = n|\theta\right) = \frac{n!}{k!\left(n-k\right)!}\theta^k\left(1-\theta\right)^{n-k}.$$

Integrating over θ, we obtain the unconditional distribution:

$$\Pr\left(\tilde{k} = k, \tilde{n} = n\right) = \frac{n!}{k!\left(n-k\right)!}\int_0^1 \theta^k\left(1-\theta\right)^{n-k} d\theta.$$

From Bayes' rule,

$$f\left(\theta|\tilde{k} = k, \tilde{n} = n\right) = \Pr\left(\tilde{k} = k, \tilde{n} = n|\theta\right) \times \frac{f\left(\theta\right)}{\Pr\left(\tilde{k} = k, \tilde{n} = n\right)}.$$

Since $\theta \sim U\left[0,1\right]$, we have

$$\begin{aligned} f\left(\theta|\tilde{k} = k, \tilde{n} = n\right) &= \frac{\Pr\left(\tilde{k} = k, \tilde{n} = n|\theta\right)}{\Pr\left(\tilde{k} = k, \tilde{n} = n\right)} = \frac{\frac{n!}{k!(n-k)!}\theta^k\left(1-\theta\right)^{n-k}}{\frac{n!}{k!(n-k)!}\int_0^1 \theta^k\left(1-\theta\right)^{n-k} d\theta} \\ &= \frac{\theta^k\left(1-\theta\right)^{n-k}}{\int_0^1 \theta^k\left(1-\theta\right)^{n-k} d\theta}. \end{aligned}$$

Moreover,

$$E\left[\theta|\tilde{k}=k,\tilde{n}=n\right]=\frac{\int_0^1\theta^{k+1}(1-\theta)^{n-k}\,d\theta}{\int_0^1\theta^{k}(1-\theta)^{n-k}\,d\theta}=\frac{k+1}{n+2}.$$

■

Lemma 8 $\int_{\frac{1}{2}}^{1}\left[\theta^k(1-\theta)^{n-k}-(1-\theta)^k\theta^{n-k}\right]d\theta\geq 0$ *if and only if* $k\geq\frac{n}{2}$.

Proof. ($\Leftarrow$) Suppose $k\geq\frac{n}{2}$. Because $n-2k\leq 0$, it follows that for all $\theta\in\left[\frac{1}{2},1\right]$, we have:

$$(1-\theta)^{n-2k}\geq\theta^{n-2k}.$$

Multiplying both sides by $[(1-\theta)\theta]^k>0$, yields

$$\theta^k(1-\theta)^{n-k}\geq(1-\theta)^k\theta^{n-k}.$$

Therefore, the term inside the integral is positive for all $\theta\in\left[\frac{1}{2},1\right]$. Integrating over this range concludes this part of the proof.

($\Rightarrow$) Suppose $k\leq\frac{n}{2}$. Then, for all $\theta\in\left[\frac{1}{2},1\right]$:

$$(1-\theta)^{n-2k}\leq\theta^{n-2k}.$$

Multiplying both sides by $[(1-\theta)\theta]^k>0$ and integrating, yields

$$\int_{\frac{1}{2}}^{1}\left[\theta^k(1-\theta)^{n-k}-(1-\theta)^k\theta^{n-k}\right]\leq 0.$$

■

In order to establish the equivalence between the equilibria of the games with the information structures presented in Figures 1 and 2, we abuse notation and let $\hat{\sigma}_H(h^t)\in\{L,H\}$ and $\hat{\sigma}_L(h^t)\in\{L,H\}$ denote the interpretations associated with a high and a low signal after history h^t. The next proposition establishes that up to a relabeling of interpretations $\hat{\sigma}_t$ there is no loss of generality in assuming that the individual always assigns a high interpretation $\hat{\sigma}=H$ to a high signal $\sigma=H$.

Note that because recollections have no intrinsic meaning, for any separating equilibrium (i.e., an equilibrium in which $\hat{\sigma}_H(h^t)\neq\hat{\sigma}_L(h^t)$), there exists an equivalent equilibrium that associates the opposite message to each signal. Moreover, for any pooling equilibrium (i.e., an equilibrium in which $\hat{\sigma}_H(h^t)=\hat{\sigma}_L(h^t)$), there exists an an equivalent equilibrium that associates the other message to both signals.

In order to deal with this uninteresting multiplicity, I will adopt the following *relabeling conditions*. Whenever we have a separating equilibrium, I will allocate each signal to its own interpretation:

$$\hat{\sigma}_H^*\left(h^t\right)\neq\hat{\sigma}_L^*\left(h^t\right)\implies\hat{\sigma}_H^*\left(h^t\right)=H\text{ and }\hat{\sigma}_L^*\left(h^t\right)=L.$$

Moreover, whenever we have a pooling equilibrium, I will allocate the high recollection to both signals:

$$\hat{\sigma}_H^*\left(h^t\right)=\hat{\sigma}_L^*\left(h^t\right)\implies\hat{\sigma}_H^*\left(h^t\right)=\hat{\sigma}_L^*\left(h^t\right)=H.$$

The following Proposition establishes the equivalence between the information structures in Figures 1 and 2. For notational clarity, I will refer to the games associated with the information structures from Figures 1 and 2 as Game 1 and Game 2.

Proposition 9 *Let $\hat{\sigma}^*$ be a profile of interpretations from Game 2 and define $(\hat{\sigma}_L^*,\hat{\sigma}_H^*)$ as $\hat{\sigma}_L^*(h^t)=\hat{\sigma}^*(h^t)$ and $\hat{\sigma}_H^*(h^t)=H$. $(\hat{\sigma}^*,a^*,\mu)$ is a PBE of Game 2 if and only if there exist beliefs μ' that coincide with μ along the equilibrium path for which $(\hat{\sigma}_L^*,\hat{\sigma}_H^*,a^*,\mu')$ is a PBE of Game 1 such that the relabeling conditions are satisfied.*

Proof. Let $(\hat{\sigma}^*, a^*, \mu)$ be a PBE of Game 2. For any history h^t, we have to establish that the individual would not profit by deviating to $\hat{\sigma} = L$ after observing $\sigma = H$ for some beliefs μ' that coincide with μ along the equilibrium path. There are two cases: $\hat{\sigma}^*(h^t) = L$ (separating equilibrium) or $\hat{\sigma}^*(h^t) = H$ (pooling equilibrium). In a separating equilibrium, deviating to $\hat{\sigma}_H = L$ would reduce the individual's self image and lead to suboptimal decisions. Therefore, deviating is not profitable.

Consider the case of a pooling equilibrium and let (off-equilibrium path) beliefs μ' attribute probability 1 to signal $\sigma_t = L$ after observing (h^t, L). Then, condition (1) from definition 1 implies that

$$\begin{aligned} &\alpha[1-\beta(1-\eta)]\sum_{s=1}^{\infty}\beta^t\left\{E\left[E_\mu\left[\theta|h^{t+s}\right]|h^t,H\right]-E\left[E_\mu\left[\theta|h^{t+s}\right]|h^t,L\right]\right\} \qquad (10)\\ \geq\;& \tilde{V}\left(a^*;\left(h^t,L\right),\left(h^t,L\right)\right)-\tilde{V}\left(a^*;\left(h^t,L\right),\left(h^t,H\right)\right). \end{aligned}$$

The individual will prefer to choose $\hat{\sigma}_H^*(h^t) = H$ in Game 1 if

$$\begin{aligned} &\alpha[1-\beta(1-\eta)]\sum_{s=1}^{\infty}\beta^t\left\{E\left[E_{\mu'}\left[\theta|h^{t+s}\right]|h^t,H\right]-E\left[E_{\mu'}\left[\theta|h^{t+s}\right]|h^t,L\right]\right\} \qquad (11)\\ \geq\;& \tilde{V}\left(a^*;\left(h^t,H\right),\left(h^t,L\right)\right)-\tilde{V}\left(a^*;\left(h^t,H\right),\left(h^t,H\right)\right). \end{aligned}$$

We have to verify that (10) implies (11). Note that $\mu(h^t,H) = \mu'(h^t,H)$ because $\hat{\sigma} = H$ is on the equilibrium path. Moreover, because μ' attributes probability 1 to signal $\sigma_t = L$ (which are the worst beliefs possible with respect to θ), it follows that $E_\mu[\theta|h^t,L] \geq E_{\mu'}[\theta|h^t,L]$. Because signals are independent, the left hand side of inequality (11) is weakly greater than the left hand side of inequality (10). Hence, a sufficient condition to ensure that (10) implies (11) is

$$\tilde{V}\left(a^*;\left(h^t,L\right),\left(h^t,L\right)\right)-\tilde{V}\left(a^*;\left(h^t,L\right),\left(h^t,H\right)\right)\geq\tilde{V}\left(a^*;\left(h^t,H\right),\left(h^t,L\right)\right)-\tilde{V}\left(a^*;\left(h^t,H\right),\left(h^t,H\right)\right). \qquad (12)$$

However, by the definition of a^*, it follows that

$$\begin{aligned} \tilde{V}\left(a^*;\left(h^t,L\right),\left(h^t,L\right)\right)-\tilde{V}\left(a^*;\left(h^t,L\right),\left(h^t,H\right)\right) &\geq 0, \text{ and}\\ \tilde{V}\left(a^*;\left(h^t,H\right),\left(h^t,L\right)\right)-\tilde{V}\left(a^*;\left(h^t,H\right),\left(h^t,H\right)\right) &\leq 0, \end{aligned}$$

which imply that (12) is satisfied.

Conversely, let $(\hat{\sigma}_L^*, \hat{\sigma}_H^*, a^*, \mu')$ be a PBE of Game 1 such that $\hat{\sigma}_L^*(h^t) = \hat{\sigma}^*(h^t)$ and $\hat{\sigma}_H^*(h^t) = H$. Then, by the definition of a PBE,

$$\hat{\sigma}_{L,t}^*\in\arg\max_{\hat{\sigma}\in\{L,H\}}\left\{\alpha[1-\beta(1-\eta)]\sum_{s=1}^{\infty}\beta^tE\left[E_{\mu'}\left[\theta|h^{t+s}\right]|h^t,\hat{\sigma}\right]+\tilde{V}\left(a^*;\left(h^t,L\right),\left(h^t,\hat{\sigma}\right)\right)\right\},\text{ and}$$

$$a_t^*\in\arg\max_{a\in A}\left\{E_{\mu'}\left[\alpha\theta+V(a,\theta)|h^t\right]\right\}.$$

Moreover, because μ' agrees with μ along the equilibrium path and the relabeling condition ensures that $\hat{\sigma}_t = H$ is on the equilibrium path for all t, it follows that the PBE conditions for Game 2 are satisfied. ■

Appendix C. Proofs

Proof of Proposition 1. For all histories h^t, let $\hat{\sigma}_t^*(h^t) = H$, $a^*(h^t) \equiv \bar{a} \in \arg\max_a \int_0^1 V(a,\theta)\,d\theta$, and $\mu(\theta|h^t) = 1$. Since the interpretation strategy does not affect beliefs and actions, it (weakly) satisfies Condition 1 from Definition 1. By construction, $a^*(h^t) = \bar{a}$ satisfies Condition 2 from Definition 1. Moreover, because there all signals are rationalized as $\hat{\sigma} = H$, consistency requires the posterior distribution to be equal to the prior distribution on the equilibrium path. Hence Condition 3 is satisfied. ■

Proof of Proposition 2. From first-order stochastic dominance, we have $\theta_t^B > \theta_{t-1}^B$ if $\sigma_t = H$ and $\theta_t^B < \theta_{t-1}^B$ if $\sigma_t = L$. Similarly, if $\hat{\sigma}_t\left(h^t, L\right) = L$, we have $\hat{\theta}_t > \hat{\theta}_{t-1}$ if $\hat{\sigma}_t = H$ and $\hat{\theta}_t < \hat{\theta}_{t-1}$ if $\sigma_t = L$. If $\hat{\sigma}_t\left(h^t, L\right) = H$, then $\hat{\theta}_t = \hat{\theta}_{t-1}$ on the equilibrium path.

Recall that θ_t^B second-order stochastically dominates $\hat{\theta}_t$ if and only if we can express θ_t^B as a mean-preserving spread of $\hat{\theta}_t$. Since signals are i.i.d., it follows that the number of high recollections k in n informative subhistories is a sufficient statistic for θ given h^t. In Appendix A (see Lemma 7), I establish that the the expected value of θ given a history with k high recollections in n informative subhistories is equal to

$$E\left[\theta|k,n\right] = \frac{k+1}{n+2}.$$

Let $\hat{\sigma}^*$ be a strategy profile such that $\hat{\sigma}_s^*\left(h^s, L\right) = H$ for some $s < t$. Consider a strategy profile $\hat{\sigma}'$ that coincides with $\hat{\sigma}^*$ except that $\hat{\sigma}_s'\left(h^s, L\right) = L$. Let θ_t' be the conditional expectation obtained by Bayes' rule given the strategy profile $\hat{\sigma}'$. Let $\epsilon_t = \theta_t' - \hat{\theta}_t$. Then,

$$\epsilon_t = \begin{cases} \frac{k+2}{n+3} - \frac{k+1}{n+2} \text{ if } \sigma_s = H \\ \frac{k+1}{n+3} - \frac{k+1}{n+2} \text{ if } \sigma_s = L \end{cases}, \text{ and}$$

$$\begin{aligned}
E\left[\epsilon_t|k,n\right] &= \Pr\left(\sigma_s = H|k,n\right)\left(\frac{k+2}{n+3} - \frac{k+1}{n+2}\right) + \Pr\left(\sigma_s = L|k,n\right)\left(\frac{k+1}{n+3} - \frac{k+1}{n+2}\right) \\
&= E\left[\theta|k,n\right]\left(\frac{k+2}{n+3} - \frac{k+1}{n+2}\right) + \left\{1 - E\left[\theta|k,n\right]\right\} \times \left(\frac{k+1}{n+3} - \frac{k+1}{n+2}\right) \\
&= \frac{k+1}{n+2}\left(\frac{k+2}{n+3} - \frac{k+1}{n+2}\right) + \left(1 - \frac{k+1}{n+2}\right) \times \left(\frac{k+1}{n+3} - \frac{k+1}{n+2}\right) = 0.
\end{aligned}$$

Thus, increasing the number of realistic interpretations leads to expected values of θ that are mean-preserving spreads of the original ones. Since θ_t^B is the expected value of θ when all interpretations are realistic, it follows that θ_t^B is a mean-preserving spread of $\hat{\theta}_t$. ■

Proof of Lemma 2. From Definition 1, the individual chooses an action a that maximizes

$$\int V\left(a,\theta\right) d\mu\left(\theta|k,n\right) = \kappa a\left[\int_{\frac{1}{2}}^{1} d\mu\left(\theta|k,n\right) - \int_0^{\frac{1}{2}} d\mu\left(\theta|k,n\right)\right].$$

Therefore, she chooses $a = 1$ if $\int_{\frac{1}{2}}^1 d\mu\left(\theta|k,n\right) \geq \int_0^{\frac{1}{2}} d\mu\left(\theta|k,n\right)$, and chooses $a = 0$ if $\int_{\frac{1}{2}}^1 d\mu\left(\theta|k,n\right) \leq \int_0^{\frac{1}{2}} d\mu\left(\theta|k,n\right)$. As shown in Appendix A (see Lemma 7), the posterior distribution has p.d.f. $\frac{\theta^k(1-\theta)^{n-k}}{\int_0^1 \theta^k(1-\theta)^{n-k} d\theta}$. Therefore, the individual chooses $a = 1$ if

$$\begin{aligned}
&\frac{\int_{\frac{1}{2}}^1 \theta^k\left(1-\theta\right)^{n-k} d\theta}{\int_0^1 \theta^k\left(1-\theta\right)^{n-k} d\theta} > \frac{\int_0^{\frac{1}{2}} \theta^k\left(1-\theta\right)^{n-k} d\theta}{\int_0^1 \theta^k\left(1-\theta\right)^{n-k} d\theta} \\
\iff &\int_{\frac{1}{2}}^1 \theta^k\left(1-\theta\right)^{n-k} d\theta > \int_0^{\frac{1}{2}} \theta^k\left(1-\theta\right)^{n-k} d\theta. \qquad (13)
\end{aligned}$$

Applying the change of variables $x = 1 - \theta$, we obtain

$$\int_0^{\frac{1}{2}} \theta^k\left(1-\theta\right)^{n-k} d\theta = \int_{\frac{1}{2}}^1 \left(1-\theta\right)^k \theta^{n-k} d\theta.$$

Plugging back inequality (13) yields

$$\int_{\frac{1}{2}}^{1}\left[\theta^{k}(1-\theta)^{n-k}-(1-\theta)^{k}\theta^{n-k}\right]d\theta>0. \tag{14}$$

As I show in the Appendix A (see Lemma 8), inequality (14) is satisfied if and only if $k \geq \frac{n}{2}$. ■

Proof of Theorem 1. Consider a history h^t with k high interpretations in n informative sub-histories. The proof proceeds by contradiction. Suppose the equilibrium strategy assigns $\hat{\sigma}(h^t)=L$ and a consider a deviation to $\hat{\sigma}_t=H$. The deviation changes beliefs to the posterior probability given $k+1$ high interpretations in n informative sub-histories, which yields a gain from self-image and an expected cost from a possibly suboptimal action.

The gain from self image in the current period is equal to $\alpha\{E[\theta|k+1,n]-E[\theta|k,n]\}=\frac{\alpha}{n+2}$. Therefore, the expected gain from self image is bounded below by $[1-\beta(1-\eta)]\times\frac{\alpha}{n+2}$, which converges to 0 at rate $\frac{1}{n}$ (this is a lower bound because we are ignoring the gains in all future periods).

Suppose the game ends after a history with $\tilde{k}$ high recollections in $\tilde{n}$ informative subhistories. Note that if $\frac{\tilde{k}}{\tilde{n}}\geq\frac{1}{2}$, the individual chooses $a=1$ both when she deviates or when she does not deviate. Similarly, the individual chooses $a=0$ in both situations when $\frac{\tilde{k}+1}{\tilde{n}}\leq\frac{1}{2}$. Therefore, deviating only generates an expected cost in terms of suboptimal actions when $\frac{\tilde{k}}{\tilde{n}}<\frac{1}{2}<\frac{\tilde{k}+1}{\tilde{n}}$. In these cases, the deviation leads the individual to choose $a=1$ instead of $a=0$, which gives an expected payoff of

$$\kappa\frac{\int_{\frac{1}{2}}^{1}\left[\theta^{\tilde{k}}(1-\theta)^{\tilde{n}-\tilde{k}}-(1-\theta)^{\tilde{k}}\theta^{\tilde{n}-\tilde{k}}\right]d\theta}{\int_0^1\theta^{\tilde{k}}(1-\theta)^{\tilde{n}-\tilde{k}}d\theta}\in(-\kappa,0).$$

Since this term is bounded below by $-\kappa$, the expected cost from suboptimal decision making is smaller than $p^*\times\kappa$, where p^* is probability of the game ending after a history such that $\frac{\tilde{k}}{\tilde{n}}<\frac{1}{2}<\frac{\tilde{k}+1}{\tilde{n}}$. In order to establish that there is a profitable deviation, it suffices to show that p^* converges to 0 faster than $\frac{1}{n}$ as $n\to\infty$ when $\frac{k}{n}$ is close to either 0 or 1.

Because the result needs to hold for any initial strategy, we need a uniform bound on p^*. First, consider an initial history h^t such that $\frac{k}{n}>\frac{1}{2}$. The probability of the game ending after a history with $\frac{\tilde{k}}{\tilde{n}}<\frac{1}{2}<\frac{\tilde{k}+1}{\tilde{n}}$ is bounded above by the strategy $\hat{\sigma}'(h)$ that interprets signals realistically until the signal becomes pivotal, and then rationalizing away every additional signal:

- $\hat{\sigma}'(h)=L$ for subsequent histories such that $k'>\frac{n'}{2}$, and
- $\hat{\sigma}'(h)=H$ for subsequent histories such that $k'\leq\frac{n'}{2}$.

Recall that $\hat{\theta}_t=\frac{k+1}{n+2}$ is the individual's expected probability of a high signal given h^t. Under the strategy above, the individual may reach a continuation history with $\frac{\tilde{k}}{\tilde{n}}<\frac{1}{2}$ if she observes:

- $2k-n$ low signals and 0 high signals, which happens with probability $\left(1-\hat{\theta}\right)^{2k-n}$,
- $2k-n+1$ low signals and 1 high signal, which happens with probability $\frac{(2k-n+2)!}{(2k-n+1)!1!}\left(1-\hat{\theta}\right)^{2k-n+1}\hat{\theta}$,
- ...
- $2k-n+s$ low signals and s high signals, which happens with probability $\frac{(2k-n+2s)!}{(2k-n+s)!s!}\left(1-\hat{\theta}\right)^{2k-n+s}\hat{\theta}^{s}$,...

Therefore, the upper bound on p^* is determined by the sum of all terms above:

$$p^*\leq\sum_{s=0}^{\infty}\frac{(2k-n+2s)!}{(2k-n+s)!s!}\left(1-\frac{k+1}{n+2}\right)^{2k-n+s}\left(\frac{k+1}{n+2}\right)^{s}.$$

Let $k = n$. Since $\lim_{n\to\infty}\left[\frac{\sum_{s=0}^{\infty}\frac{(n+2s)!}{(n+s)!s!}\left(1-\frac{n+1}{n+2}\right)^{n+s}\left(\frac{n+1}{n+2}\right)^{s}}{\frac{1}{n}}\right] = 0$, it follows that for large n

$$[1-\beta(1-\eta)]\times\frac{\alpha}{n+2} > \kappa\times\sum_{s=0}^{\infty}\frac{(n+2s)!}{(n+s)!s!}\left(1-\frac{n+1}{n+2}\right)^{n+s}\left(\frac{n+1}{n+2}\right)^{s} \geq p^*\kappa.$$

Therefore, by continuity, there exist $\gamma_H < 1$ and $\bar{n}$ such that whenever $\frac{k}{n} > \gamma_H$ and $n > \bar{n}$, $\frac{\alpha}{n+2} > p^*\kappa$. This establishes that it is profitable to deviate to $\hat{\sigma}_{t+1} = H$ in this case, which concludes the first part of the proof.

Now consider an initial history with $\frac{k+1}{n} < \frac{1}{2}$. Again, a uniform bound on p^* is attained by playing the following strategy:

- $\hat{\sigma}(h) = L$ for subsequent histories such that $k'+1 < \frac{n'}{2}$, and
- $\hat{\sigma}(h) = H$ for all subsequent histories such that $k'+1 \geq \frac{n'}{2}$.

Under the strategy above, the individual may reach a continuation history with $k'+1 \geq \frac{n'}{2}$ if she observes $n-2k+f-2$ high signals and f low signals, which happens with probability $\frac{(n-2k+2f-2)!}{(n-2k+f-2)!f!} \times \left(1-\hat{\theta}\right)^{f}\hat{\theta}^{n-2k+f-2}$, for any $f \in \mathbb{N}$. Thus, the probability of the signal being pivotal when the game ends under this strategy is equal to:

$$\kappa\sum_{f=0}^{\infty}\frac{(n-2k+2f-2)!}{(n-2k+f-2)!f!}\times\left(1-\frac{k+1}{n+2}\right)^{f}\left(\frac{k+1}{n+2}\right)^{n-2k+f-2}.$$

Take $k = 0$. Using the fact that $\lim_{n\to\infty}\kappa\frac{\sum_{f=0}^{\infty}\frac{(n+2f-2)!}{(n+f-2)!f!}\times\left(1-\frac{1}{n+2}\right)^{f}\left(\frac{1}{n+2}\right)^{n+f-2}}{\frac{1}{n}} = 0$, it follows that there exist $\gamma_L > 0$ and $\bar{n}$ such that whenever $\frac{k}{n} < \gamma_L$ and $n > \bar{n}$,

$$[1-\beta(1-\eta)]\times\frac{\alpha}{n+2} > \kappa\sum_{f=0}^{\infty}\frac{(n-2k+2f-2)!}{(n-2k+f-2)!f!}\left(1-\frac{k+1}{n+2}\right)^{f}\left(\frac{k+1}{n+2}\right)^{n-2k+f-2} \geq p^*\kappa,$$

which concludes the proof. ■

Proof of Proposition 3. Suppose $\hat{\sigma}^*_M(0,0) = L$ and consider a deviation to H. The increase in self-image in a period associated with state (k,n) is equal to

$$E[\theta|k+1,n] - E[\theta|k,n] = \frac{1}{n+2}. \tag{15}$$

We need to obtain a uniform lower bound on the discounted benefit from deviating. Since the expression in (15) is decreasing in the number of informative subhistories, the discounted gain in self-image is bounded below by following the strategy that interprets every additional signal realistically, which yields

$$\begin{aligned}\alpha[1-\beta(1-\eta)]\sum_{t=0}^{+\infty}\frac{\beta^t(1-\eta)^t}{t+3} &= \alpha[1-\beta(1-\eta)]\left(\frac{1}{3}+\sum_{t=1}^{+\infty}\frac{\beta^t(1-\eta)^t}{t+3}\right)\\ &\geq \frac{\alpha}{3}[1-\beta(1-\eta)].\end{aligned}$$

We want to obtain a uniform upper bound on the cost of taking a suboptimal action. When the final state is such that either $\frac{\tilde{k}}{\tilde{n}} > \frac{1}{2}$ or $\frac{\tilde{k}+1}{\tilde{n}} < \frac{1}{2}$, the deviation does not affect actions. Thus, the cost is positive only states such that $\frac{\tilde{k}}{\tilde{n}} < \frac{1}{2} < \frac{\tilde{k}+1}{\tilde{n}}$. In these cases, rationalization of the signal induces the individual to take action $a = 1$ when $a = 0$ is preferred, and the cost of taking the suboptimal action is equal to

$$-\kappa \frac{\int_{\frac{1}{2}}^{1} \left[\theta^{\tilde{k}} (1-\theta)^{\tilde{n}-\tilde{k}} - (1-\theta)^{\tilde{k}} \theta^{\tilde{n}-\tilde{k}}\right] d\theta}{\int_0^1 \theta^{\tilde{k}} (1-\theta)^{\tilde{n}-\tilde{k}} d\theta} \in (-\kappa, 0),$$

where $\frac{\tilde{k}}{\tilde{n}} < \frac{1}{2} < \frac{\tilde{k}+1}{\tilde{n}}$. Thus, the cost is bounded above by κ. Since $\kappa < \frac{\alpha}{3} [1 - \beta(1-\eta)]$, the individual has a profitable deviation. ■

Proof of Proposition 4. If the individual observes a low signal at state $(0,0)$, she interprets it realistically and moves to (absorbing) state $(0,1)$. Therefore, she obtains an expected payoff of

$$\alpha [1 - \beta(1-\eta)] \left(1 + \beta(1-\eta) + \beta^2 (1-\eta)^2 + \beta^3 (1-\eta)^3 + ..\right) \frac{1}{3} + 0 = \frac{\alpha}{3}$$

(because $a^*_M(0,1) = 0$). By deviating to $\hat{\sigma} = H$, the individual moves to (absorbing) state $(1,1)$ and gets an expected payoff of

$$\frac{2\alpha}{3} + \kappa \frac{\int_{\frac{1}{2}}^{1} [(1-\theta) - \theta] d\theta}{\int_0^1 (1-\theta) d\theta} = \frac{2\alpha}{3} - \frac{\kappa}{2}.$$

This deviation is not profitable if $\frac{\alpha}{3} \geq \frac{2\alpha}{3} - \frac{\kappa}{2}$, which is equivalent to $\kappa \geq \frac{2\alpha}{3}$.

Let (off-equilibrium path) beliefs be such that, upon observing $\hat{\sigma} = L$ at any state, the individual assigns probability 1 to signal $\sigma = L$. Then, deviating to $\hat{\sigma} = L$ at state $(0,1)$ leads to a payoff of

$$\alpha E[\theta | (0,2)] + E[V(a(0,2), \theta) | 0, 2] = \frac{\alpha}{4},$$

which is lower than the payoff from following the equilibrium strategy and playing $\hat{\sigma} = H$ because it lowers the payoff from self-image but does not affect the payoff from actions. Therefore, this is not a profitable deviation.

At (absorbing) state $(1,1)$, the individual plays $\hat{\sigma} = H$ and obtains an expected payoff of

$$\alpha E[\theta | (1,1)] + \underbrace{E[V(a(1,2), \theta) | 1, 2]}_{0} = \frac{2\alpha}{3}.$$

By deviating to $\hat{\sigma} = L$, she obtains $E[\alpha\theta | (1,2)] + \underbrace{E[V(a(1,2), \theta) | 1, 2]}_{0} = \frac{\alpha}{2}$, which is lower than the equilibrium payoff. Thus, it is not profitable to deviate at state $(1,1)$. ■

Proof of Theorem 2. Fix a type $\theta \notin \left\{0, \frac{1}{2}, 1\right\}$. In order to obtain a contradiction, suppose that the claim above is not true. Recall that (k,n) is an absorbing state if $\hat{\sigma}^*_M(k,n) = H$. Therefore, there must exist a sequence $\{k_n\}_{n=0}^{\infty}$ such that

$$k_0 = 0, \quad k_{n+1} \in \{k_n, k_n + 1\},$$

and $\hat{\sigma}^*_M(k_n, n) = L$ for all $n \in \mathbb{N}$.

Define the random variable $x_t \equiv \mathbf{1}(\hat{\sigma}_t = H)$, where $\mathbf{1}(.)$ denotes the indicator function. Then, each x_t is an independent Bernoulli(θ) random variable. Denote the sum of x_t by $S_n \equiv X_1 + ... + X_n$, and define the variable

$$Z_n \equiv \sqrt{n} \times \frac{\frac{S_n}{n} - \theta}{\sqrt{\theta(1-\theta)}}.$$

Then, by the Central Limit Theorem, $Z_n \to_d N(0,1)$.

As in the proof of Theorem 1, deviating from the equilibrium choice of $\hat{\sigma}^*_M(k_n, n) = L$ and playing $\hat{\sigma} = H$ leads to a gain in self-image and a loss due to worse decision making. The gain from self image is bounded below by

$$\alpha\left[1-\beta\left(1-\eta\right)\right]\left\{E\left[\theta|k+1,n\right]-E\left[\theta|k,n\right]\right\}=\frac{\alpha}{n+2}\left[1-\beta\left(1-\eta\right)\right],$$

which converges to $\frac{\alpha(1-\beta)}{n+2}$ as $\eta \to 0$.

The decision making loss is bounded above by $\kappa \Pr\left(\frac{S_n}{n} < \frac{1}{2} < \frac{S_n+1}{n}\right)$. Rewriting in terms of Z_n, yields:

$$\kappa \Pr\left(Z_n < \sqrt{n}\frac{\frac{1}{2}-\theta}{\sqrt{\theta\left(1-\theta\right)}} < Z_n + \frac{1}{\sqrt{n}\sqrt{\theta\left(1-\theta\right)}}\right).$$

Because $Z_n \to_d N\left(0,1\right)$, it follows that

$$\Pr\left(\frac{S_n}{n} < \frac{1}{2} < \frac{S_n+1}{n}\right) \approx \Phi\left(\sqrt{n}\frac{\frac{1}{2}-\theta}{\sqrt{\theta\left(1-\theta\right)}}\right) - \Phi\left(\sqrt{n}\frac{\frac{1}{2}-\theta}{\sqrt{\theta\left(1-\theta\right)}} - \frac{1}{\sqrt{n}\sqrt{\theta\left(1-\theta\right)}}\right),$$

where Φ is the c.d.f. of a standard normal variable.

In order to establish the desired contradiction, we need to show that the gain from self image is greater than the decision making cost:

$$\frac{\alpha\left(1-\beta\right)}{\kappa\left(n+2\right)} > \Phi\left(\sqrt{n}\frac{\frac{1}{2}-\theta}{\sqrt{\theta\left(1-\theta\right)}}\right) - \Phi\left(\sqrt{n}\frac{\frac{1}{2}-\theta}{\sqrt{\theta\left(1-\theta\right)}} - \frac{1}{\sqrt{n}\sqrt{\theta\left(1-\theta\right)}}\right) \tag{16}$$

for n sufficiently large.

Note that both sides of inequality (16) converge to 0 as $n \to \infty$. Using the formula for Φ, it follows that $\Phi\left(z\right) - \Phi\left(y\right) = \frac{1}{\sqrt{2\pi}}\int_y^z \exp\left(-\frac{x^2}{2}\right)dx$. Applying a Taylor approximation, we obtain

$$\Phi\left(y\right) \approx \Phi\left(z\right) + \left(y-z\right)\phi\left(y\right),$$

where $\phi\left(x\right) = \frac{1}{\sqrt{2\pi}}\exp\left(-\frac{x^2}{2}\right)$. Then, we have

$$\begin{aligned}&\Phi\left(\sqrt{n}\frac{\frac{1}{2}-\theta}{\sqrt{\theta\left(1-\theta\right)}}\right) - \Phi\left(\sqrt{n}\frac{\frac{1}{2}-\theta}{\sqrt{\theta\left(1-\theta\right)}} - \frac{1}{\sqrt{n}\sqrt{\theta\left(1-\theta\right)}}\right)\\ \approx\ & \frac{1}{\sqrt{2\pi}\sqrt{n}\sqrt{\theta\left(1-\theta\right)}}\exp\left[-\frac{n\left(\frac{\frac{1}{2}-\theta}{\sqrt{\theta(1-\theta)}}\right)^2}{2}\right].\end{aligned}$$

Then, in order to show that inequality (16) is satisfied for large n, it suffices to show that

$$\frac{\alpha\left(1-\beta\right)}{\kappa\left(n+2\right)} > \frac{1}{\sqrt{2\pi}\sqrt{n}\sqrt{\theta\left(1-\theta\right)}}\exp\left[-\frac{n\left(\frac{\frac{1}{2}-\theta}{\sqrt{\theta(1-\theta)}}\right)^2}{2}\right].$$

Rearranging, we obtain $\frac{\alpha(1-\beta)}{\kappa}\frac{\sqrt{n}}{n+2} > \frac{1}{\sqrt{2\pi}\sqrt{\theta(1-\theta)}}\exp\left[-\frac{n\left(\frac{\frac{1}{2}-\theta}{\sqrt{\theta(1-\theta)}}\right)^2}{2}\right]$. Since $\frac{\sqrt{n}}{n+2} < \frac{1}{\sqrt{n}}$, a sufficient condition is $\frac{\alpha(1-\beta)}{\kappa}\frac{1}{\sqrt{n}} > \frac{1}{\sqrt{2\pi}\sqrt{\theta(1-\theta)}}\exp\left[-\frac{n\left(\frac{\frac{1}{2}-\theta}{\sqrt{\theta(1-\theta)}}\right)^2}{2}\right]$.Rearranging, yields

$$\frac{\alpha(1-\beta)\sqrt{2\pi}}{\kappa} > \frac{\sqrt{n}}{\sqrt{\theta(1-\theta)}}\exp\left[-n\frac{(1-2\theta)^2}{8\theta(1-\theta)}\right].$$

Note that the term on the left is a positive constant whereas the term on the right is a function of n. Moreover, for $\theta \neq \frac{1}{2}$, we have $\lim_{n\to\infty}\sqrt{n}\exp\left[-n\frac{(1-2\theta)^2}{8\theta(1-\theta)}\right] = 0$. Thus, it follows that for large n,

$$\frac{\sqrt{n}}{\sqrt{\theta(1-\theta)}}\exp\left[-n\frac{(1-2\theta)^2}{8\theta(1-\theta)}\right] < \frac{\alpha(1-\beta)\sqrt{2\pi}}{\kappa},$$

which concludes the argument for $\theta \notin \left\{0, \frac{1}{2}, 1\right\}$.

Now suppose $\theta = 0$ (the case where $\theta = 1$ is analogous). Then, we must have $\sigma_t = L$ for all t. In order to obtain a contradiction, suppose that the claim is not true. Then, because of the Markovian restriction, we must have $\hat{\sigma}_t^* = L$ for all t. The individual's gain from deviating to $\hat{\sigma} = H$ is bounded below by

$$\alpha[1-\beta(1-\eta)]\left(\frac{2}{n+2} - \frac{1}{n+2}\right) = \frac{\alpha[1-\beta(1-\eta)]}{n+2} > 0.$$

For $n > 2$, the deviation leads the individual to take action $a(1,n) = 0$ since $\frac{1}{n} < \frac{1}{2}$. Therefore, deviating to $\hat{\sigma} = H$ is profitable since it increases the individual's self views but does not affect her actions. ■

Proof of Theorem 3. In order to obtain a contradiction, suppose that $\hat{\sigma}(h^t) = L$ for some history h^t with n large enough and consider a deviation to $\hat{\sigma} = H$. The gain from self-image is bounded below by

$$\alpha[1-\beta(1-\eta)]\{E[\theta|k+1,n] - E[\theta|k,n]\} = \frac{\alpha[1-\beta(1-\eta)]}{n+2}.$$

Let $\tilde{\gamma} \equiv \frac{\tilde{k}}{\tilde{n}}$, and $\tilde{\gamma}' \equiv \frac{\tilde{k}+1}{\tilde{n}}$. The deviation will lead the individual to choose action

$$a(\tilde{\gamma}', \tilde{n}) = \arg\max_{a'}\int V(a',\theta)\, d\mu(\theta|\tilde{n}\tilde{\gamma}', \tilde{n})$$

when the game ends after a history associated with $\left(\tilde{k}, \tilde{n}\right) = (\tilde{n}\tilde{\gamma}, \tilde{n})$. Therefore, the expected cost of distorting actions is equal to

$$\sum_{\substack{\tilde{n}\tilde{\gamma}\geq k\\ \tilde{n}\geq n}}\left[\int V\left(a\left(\tilde{\gamma}', \tilde{n}\right), \theta\right) d\mu(\theta|\tilde{n}\tilde{\gamma}, \tilde{n}) - \int V(a(\tilde{\gamma}, \tilde{n}), \theta)\, d\mu(\theta|\tilde{n}\tilde{\gamma}, \tilde{n})\right]\Phi(\tilde{n}\tilde{\gamma}, \tilde{n}|k,n), \qquad (17)$$

where $\Phi\left(\tilde{k}, \tilde{n}|k,n\right)$ denotes the probability that the game ends at a history associated with $\left(\tilde{k}, \tilde{n}\right)$ consistent with the strategies being played in the fixed PBE. We will establish that the expression in (17) decreases to 0 at a rate faster than $\frac{1}{n}$.

For notational simplicity, I will omit the term $\tilde{n}$ from $a(\tilde{\gamma}, \tilde{n})$. Define ξ as

$$\xi(a(\tilde{\gamma}), \tilde{\gamma}, \tilde{n}) \equiv \int V(a(\tilde{\gamma}, \tilde{n}), \theta)\, d\mu(\theta|\tilde{n}\tilde{\gamma}, \tilde{n}).$$

Then, from Taylor's theorem with the Lagrange remainder, it follows that

$$\xi(a(\tilde{\gamma}), \tilde{\gamma}, \tilde{n}) = \xi\left(a\left(\tilde{\gamma}'\right), \tilde{\gamma}, \tilde{n}\right) + \frac{\partial\xi}{\partial a}(a(\tilde{\gamma}), \tilde{\gamma}, \tilde{n}) \times \left[a\left(\tilde{\gamma}'\right) - a(\tilde{\gamma})\right] + r\left(a\left(\tilde{\gamma}'\right) - a(\tilde{\gamma})\right), \qquad (18)$$

where $\lim_{h\to 0}\frac{r(h)}{h}=0$. Since $a(\hat{\gamma})\in\max_a\int V(a,\theta)\,d\mu(\theta|\tilde{n}\tilde{\gamma},\tilde{n})$, it must satisfy the following local first-order condition:

$$\frac{\partial\xi}{\partial a}(a(\tilde{\gamma}),\tilde{\gamma},\tilde{n})=0.$$

Therefore, equation (18) becomes

$$\xi(a(\tilde{\gamma}),\tilde{\gamma},\tilde{n})=\xi\left(a\left(\tilde{\gamma}'\right),\tilde{\gamma},\tilde{n}\right)+r\left(a\left(\tilde{\gamma}'\right)-a(\tilde{\gamma})\right).$$

It remains to be shown that $\lim_{n\to\infty}\frac{r\left(a\left(\tilde{\gamma}'\right)-a(\tilde{\gamma})\right)}{\frac{1}{n}}=0$. Note that $\frac{1}{n}=\tilde{\gamma}'-\tilde{\gamma}$, so that

$$\begin{aligned}\frac{r\left(a\left(\tilde{\gamma}'\right)-a(\tilde{\gamma})\right)}{\frac{1}{n}}&=\frac{r\left(a\left(\tilde{\gamma}'\right)-a(\tilde{\gamma})\right)}{\tilde{\gamma}'-\tilde{\gamma}}\\&=\frac{r\left(a\left(\tilde{\gamma}'\right)-a(\tilde{\gamma})\right)}{a\left(\tilde{\gamma}'\right)-a(\tilde{\gamma})}\frac{a\left(\tilde{\gamma}'\right)-a(\tilde{\gamma})}{\tilde{\gamma}'-\tilde{\gamma}}.\end{aligned}$$

Moreover, from the definition of r, we have $\lim_{a(\tilde{\gamma}')-a(\tilde{\gamma})\to 0}\frac{r\left(a\left(\tilde{\gamma}'\right)-a(\tilde{\gamma})\right)}{a(\tilde{\gamma}')-a(\tilde{\gamma})}=0$. Thus, it suffices to show that there exists some constants ε and $K_{\tilde{\gamma}}$ such that $\left|\tilde{\gamma}'-\tilde{\gamma}\right|<\varepsilon$ implies

$$\left|\frac{a\left(\tilde{\gamma}'\right)-a(\tilde{\gamma})}{\tilde{\gamma}'-\tilde{\gamma}}\right|<K_{\tilde{\gamma}}.$$

By the Theorem of the Maximum, a is continuous. The result then follows from the fact that continuous functions have bounded subdifferentials at any interior points.

Therefore, for n large enough, there exists some $n^*(k,\eta)$ such that $\hat{\sigma}(h^t)=H$ for all $n>n^*(k,\eta)$, $k\in\{0,...,n\}$. Since $\{0,...,n\}$ is a finite set, letting $\bar{n}(\eta)\equiv\max_{k\in\{0,...,n\}}n^*(k,\eta)$ concludes the proof. ■

Proof of Lemma 3. Let $X_1,X_2,...,X_n$ be a sequence of independent Bernoulli random variables with parameter $\pi(\theta)$. Therefore, they are distributed according to a probability mass function

$$f(x|\theta)=\pi(\theta)^x\,[1-\pi(\theta)]^{1-x},\ x\in\{0,1\}.$$

It is straightforward to check that Assumptions 1-9 from Johnson (1970) are satisfied. Moreover,

$$\int_{\underline{\theta}}^{\bar{\theta}}|\theta|\,\rho(\theta)\,d\theta\le\bar{\theta}<\infty.$$

Therefore, from the Theorem 3.1 in Johnson (1970), it follows that there exists constants C and N_k such that

$$\left|E[\theta|X_1,X_2,...,X_n]-\hat{\theta}-\frac{1}{b\left(\hat{\theta}\right)}\left(6a_3\left(\hat{\theta}\right)+\frac{\rho'\left(\hat{\theta}\right)}{\rho\left(\hat{\theta}\right)}\right)n^{-1}\right|\le Cn^{-2},\tag{19}$$

where $\hat{\theta}$ is the maximum likelihood estimator of θ,

$$\begin{aligned}a_{2n}(\theta)&\equiv\frac{1}{2n}\sum_{i=1}^{n}\frac{\partial^2}{\partial\theta^2}\log f(x_i|\theta),\qquad(20)\\a_{3n}(\theta)&\equiv\frac{1}{6n}\sum_{i=1}^{n}\frac{\partial^3}{\partial\theta^3}\log f(x_i|\theta),\ \text{and}\\b(\theta)&\equiv\sqrt{-2a_{2n}(\theta)}.\end{aligned}$$

Note that the maximum likelihood estimator of $\pi(\theta)$ in this model is $\widehat{\pi(\theta)} = \frac{\sum_{i=1}^{n} x_i}{n}$. Moreover, from the invariance property, it follows that $\hat{\theta} = \pi^{-1}\left(\frac{\sum_{i=1}^{n} x_i}{n}\right)$. Computing the terms in equations (20) and substituting in (19), we obtain:

$$\left| \begin{array}{c} E[\theta|X_1, X_2, ..., X_n] - \hat{\theta} \\ -\frac{\sqrt{\pi(\hat{\theta})[1-\pi(\hat{\theta})]}}{\pi'(\hat{\theta})}\left(\frac{1}{\pi(\hat{\theta})[1-\pi(\hat{\theta})]}\left\{2\left[\pi'\left(\hat{\theta}\right)\right]^3 \frac{1-2\pi(\hat{\theta})}{\pi(\hat{\theta})[1-\pi(\hat{\theta})]} - 3\pi'\left(\pi^{-1}\left(\hat{\theta}\right)\right)\pi''\left(\hat{\theta}\right)\right\} + \frac{\rho'(\hat{\theta})}{\rho(\hat{\theta})}\right)n^{-1} \end{array} \right| \tag{21}$$
$$\leq \quad Cn^{-2}.$$

As before, let $k \equiv \sum_{i=1}^{n} x_i$ denote the number of successes in the n Bernoulli trials. Then, substituting $\hat{\theta} = \pi^{-1}\left(\frac{\sum_{i=1}^{n} x_i}{n}\right)$ yields

$$\left|E[\theta|X_1, X_2, ..., X_n] - \pi^{-1}\left(\frac{k}{n}\right) - \xi\left(\frac{k}{n}\right)n^{-1}\right| \leq Cn^{-2},$$

where ξ is given by equation (4). ■

Proof of Lemma 4. From equation (3), there exist C_0 and C_1 such that

$$\left|E[\theta|k, n] - \pi^{-1}\left(\frac{k}{n}\right) - \xi\left(\frac{k}{n}\right)n^{-1}\right| \leq C_0 n^{-2}, \text{ and}$$

$$\left|E[\theta|k+1, n] - \pi^{-1}\left(\frac{k+1}{n}\right) - \xi\left(\frac{k+1}{n}\right)n^{-1}\right| \leq C_1 n^{-2}$$

for all $n > \max\{N_k, N_{k+1}\}$ on an almost sure set (under the true θ). From the triangular inequality, it follows that

$$\left|E[\theta|k+1, n] - E[\theta|k, n] - \pi^{-1}\left(\frac{k+1}{n}\right) + \pi^{-1}\left(\frac{k}{n}\right) - \xi\left(\frac{k+1}{n}\right)n^{-1} + \xi\left(\frac{k}{n}\right)n^{-1}\right| \leq C_2 n^{-2} \tag{22}$$

for $C_2 \equiv C_0 + C_1$.

Let $\hat{\theta}^* \equiv \pi^{-1}\left(\frac{k+1}{n}\right)$ and recall that $\hat{\theta} = \frac{k}{n}$. Applying a Taylor approximation, it follows that there exist constants Z_0 and $N_{1,k}$ such that

$$\left|\xi\left(\hat{\theta}^*\right) - \xi\left(\hat{\theta}\right) - \left(\hat{\theta}^* - \hat{\theta}\right)\xi'\left(\hat{\theta}\right)\right| \leq Z_0 n^{-2}$$

for all $n > N_{1,k}$. Substituting in equation (22), we obtain

$$\left|E[\theta|k+1, n] - E[\theta|k, n] - \pi^{-1}\left(\frac{k+1}{n}\right) + \pi^{-1}\left(\frac{k}{n}\right) - \left[\pi^{-1}\left(\frac{k+1}{n}\right) - \pi^{-1}\left(\frac{k}{n}\right)\right]\xi'\left(\frac{k}{n}\right)n^{-1}\right| \tag{23}$$
$$\leq \quad Z_1 n^{-2}$$

for constants $N_{2,k}$ and Z_1 and $n > N_{2,k}$ on an almost sure set.

Since π is twice continuously differentiable, there exist constants $N_{3,k}$ and Z_2 such that for all $n > N_{3,k}$ and Z_2

$$\left|\pi^{-1}\left(\frac{k+1}{n}\right) - \pi^{-1}\left(\frac{k}{n}\right) - \frac{1}{n}\frac{1}{\pi'\left(\pi^{-1}\left(\frac{k}{n}\right)\right)}\right| < Z_2 n^{-2}.$$

Then, equation (23) implies that there exist constants $N_{4,k}$ and Z_3 such that for all $n > N_{4,k}$:

$$\left|E[\theta|k+1, n] - E[\theta|k, n] - \frac{1}{\pi'\left(\pi^{-1}\left(\frac{k}{n}\right)\right)}n^{-1} - \frac{\xi'\left(\frac{k}{n}\right)}{\pi'\left(\pi^{-1}\left(\frac{k}{n}\right)\right)}n^{-2}\right| \leq Z_3 n^{-2}$$

on an almost sure set. Letting $D_1 \equiv Z_3 + \left|\frac{\xi'\left(\frac{k}{n}\right)}{\pi'\left(\pi^{-1}\left(\frac{k}{n}\right)\right)}\right|$, this inequality implies that

$$\left|E\left[\theta|k+1,n\right] - E\left[\theta|k,n\right] - \frac{1}{\pi'\left(\pi^{-1}\left(\frac{k}{n}\right)\right)} n^{-1}\right| \leq D_1 n^{-2},$$

which concludes the proof. ■

Proof of Theorem 4. The proof follows the exact same argument as the proof of Theorem **3.** ■

Theorem 5 will be established by a series of lemmata:

Lemma 9 $\sqrt{n}\left\{E\left[\pi|k,n\right] - \pi\right\} \to_D N\left(0, \frac{1}{b^2}\right)$, *where* $b = \sqrt{-\frac{1}{n}\left\{\sum_{i=1}^{n} \frac{\partial^2}{\partial \pi^2} \log f\left(x_i|\pi\right)\right\}_{\pi=\hat{\pi}}} = \frac{1}{\sqrt{\hat{\pi}(1-\hat{\pi})}}$, *and* $\hat{\pi} = \frac{k}{n}$ *is the MLE of* π.

Proof. This is a straightforward application of the Laplace-Bernstein-von Mises theorem (see Theorem 1.4.3 and Remark 1.4.6 in Ghosh and Ramamoorthi, 2003). ■

Lemma 10 *Suppose* $\rho\left(.\right)$ *is regular. Then, there exists constants* D *and* N_k *such that for any* $n > N_k$:

$$\left|E\left[\pi|k+1,n\right] - E\left[\pi|k,n\right] - \frac{1}{n}\right| \leq D\frac{1}{n^2}.$$

Proof. Let p denote the prior distribution over π implied by $\rho : p\left(\pi\left(\theta\right)\right) = \rho\left(\theta\right)$. Note that

$$p'\left(\pi\left(\theta\right)\right)\pi'\left(\theta\right) = \rho'\left(\theta\right) \therefore p'\left(\pi\left(\theta\right)\right) = \frac{\rho'\left(\theta\right)}{\pi'\left(\theta\right)}.$$

Thus, $p'\left(\pi\right) = \frac{\rho'\left(\theta^{-1}(\pi)\right)}{\pi'\left(\theta^{-1}(\pi)\right)}$ where $\theta^{-1}\left(\pi\right)$ is the inverse of the relation $\pi\left(\theta\right)$, i.e. $\pi\left(\theta^{-1}\left(\pi\right)\right) = \pi$. From Johnson (1970), there exist constants D_1 and N such that $n > N$ implies

$$\left|E_X\left[\pi\right] - \hat{\pi} - \sqrt{\hat{\pi}\left(1-\hat{\pi}\right)}\left[2\frac{1-2\hat{\pi}}{\hat{\pi}^2\left(1-\hat{\pi}\right)^2} + \frac{p'\left(\hat{\pi}\right)}{p\left(\hat{\pi}\right)}\right] n^{-1}\right| \leq D_1 n^{-2}.$$

where E_X denotes the posterior distribution of π given some sequence of signals X and $\hat{\pi}$ denotes the maximum likelihood estimator of π. Since $p'\left(\hat{\pi}\right) = \frac{\rho'\left(\theta^{-1}(\hat{\pi})\right)}{\pi'\left(\theta^{-1}(\hat{\pi})\right)}$ and $\hat{\pi} = \frac{k}{n}$, we obtain

$$\left|E_X\left[\pi\right] - \frac{k}{n} - \sqrt{\frac{k}{n}\left(1-\frac{k}{n}\right)}\left[2\frac{1-2\frac{k}{n}}{\left(\frac{k}{n}\right)^2\left(1-\frac{k}{n}\right)^2} + \frac{\rho'\left(\theta^{-1}\left(\frac{k}{n}\right)\right)}{p\left(\frac{k}{n}\right)\pi'\left(\theta^{-1}\left(\frac{k}{n}\right)\right)}\right] n^{-1}\right| \leq D_1 n^{-2}.$$

Let $\xi\left(\hat{\pi}\right) \equiv \sqrt{\hat{\pi}\left(1-\hat{\pi}\right)}\left[2\frac{1-2\hat{\pi}}{\hat{\pi}^2(1-\hat{\pi})^2} + \frac{\rho'\left(\theta^{-1}(\hat{\pi})\right)}{p(\hat{\pi})\pi'\left(\theta^{-1}(\hat{\pi})\right)}\right]$. Then,

$$E\left[\pi|k,n\right] = \frac{k}{n} + \xi\left(\frac{k}{n}\right) n^{-1} + O\left(n^{-2}\right), \tag{24}$$

Applying a Taylor expansion yields

$$\xi\left(\frac{k+1}{n}\right) = \xi\left(\frac{k}{n}\right) + \frac{1}{n}\xi'\left(\frac{k}{n}\right) + O\left(\frac{1}{n^2}\right).$$

Therefore,

$$E[\pi|k+1,n] = \frac{k+1}{n} + \xi\left(\frac{k}{n}\right)n^{-1} + O\left(n^{-2}\right) \tag{25}$$

Combining (24) and (25) yields

$$\begin{aligned} E[\pi|k+1,n] - E[\pi|k,n] &= \frac{k+1}{n} + \xi\left(\frac{k}{n}\right)n^{-1} - \frac{k}{n} - \xi\left(\frac{k}{n}\right)n^{-1} + O\left(n^{-2}\right) \\ &= \frac{1}{n} + O\left(n^{-2}\right), \end{aligned}$$

which concludes the proof. ■

Proof of Theorem 5. Under Assumption 3, a signal is pivotal if

$$E[\pi|k,n] < \lambda < E[\pi|k+1,n]. \tag{26}$$

Applying the result from Lemma 10, we obtain

$$E[\pi|k+1,n] < E[\pi|k,n] + \frac{1}{n} + Dn^{-2}$$

for $n > N_k$. Therefore, inequality (26) becomes

$$\frac{\sqrt{n}}{\sqrt{\frac{k}{n}\left(1-\frac{k}{n}\right)}}(\lambda-\pi) - \frac{\sqrt{n}\left(Cn^{-2}+\frac{1}{n}\right)}{\sqrt{\frac{k}{n}\left(1-\frac{k}{n}\right)}} < \frac{\sqrt{n}}{\sqrt{\frac{k}{n}\left(1-\frac{k}{n}\right)}}\left\{E[\pi|k,n]-\pi\right\} < \frac{\sqrt{n}}{\sqrt{\frac{k}{n}\left(1-\frac{k}{n}\right)}}(\lambda-\pi)$$

for some constant C.

Since $\frac{\sqrt{n}}{\sqrt{\frac{k}{n}\left(1-\frac{k}{n}\right)}}\left\{E[\pi|k,n]-\pi\right\} \to_D N(0,1)$, the probability of a signal being pivotal is bounded above by the following probability:

$$\Phi\left(\frac{\sqrt{n}}{\sqrt{\frac{k}{n}\left(1-\frac{k}{n}\right)}}(\lambda-\pi)\right) - \Phi\left(\frac{\sqrt{n}}{\sqrt{\frac{k}{n}\left(1-\frac{k}{n}\right)}}(\lambda-\pi) - \frac{\sqrt{n}\left(Cn^{-2}+\frac{1}{n}\right)}{\sqrt{\frac{k}{n}\left(1-\frac{k}{n}\right)}}\right).$$

Note that this term converges to zero. We want to show that

$$\frac{\Phi\left(\frac{\sqrt{n}}{\sqrt{\frac{k}{n}\left(1-\frac{k}{n}\right)}}(\lambda-\pi)\right) - \Phi\left(\frac{\sqrt{n}}{\sqrt{\frac{k}{n}\left(1-\frac{k}{n}\right)}}(\lambda-\pi) - \frac{\sqrt{n}\left(Cn^{-2}+\frac{1}{n}\right)}{\sqrt{\frac{k}{n}\left(1-\frac{k}{n}\right)}}\right)}{\frac{1}{n}} \to_{a.s.} 0,$$

which will establish that the gain from self-perception is greater than the cost of suboptimal decisions when n is high for almost all histories.

From Taylor's theorem, we obtain

$$\begin{aligned} &\Phi\left(\frac{\sqrt{n}}{\sqrt{\frac{k}{n}\left(1-\frac{k}{n}\right)}}(\lambda-\pi)\right) - \Phi\left(\frac{\sqrt{n}}{\sqrt{\frac{k}{n}\left(1-\frac{k}{n}\right)}}(\lambda-\pi) - \frac{\sqrt{n}\left(Cn^{-2}+\frac{1}{n}\right)}{\sqrt{\frac{k}{n}\left(1-\frac{k}{n}\right)}}\right) \\ = &\frac{\sqrt{n}\left(Cn^{-2}+\frac{1}{n}\right)}{\sqrt{\frac{k}{n}\left(1-\frac{k}{n}\right)}}\frac{1}{\sqrt{2\pi}}\exp\left\{-\frac{\left[\frac{\sqrt{n}}{\sqrt{\frac{k}{n}\left(1-\frac{k}{n}\right)}}(\lambda-\pi)\right]^2}{2}\right\} + O\left(\frac{1}{n^2}\right). \end{aligned}$$

Therefore, it suffices that establish that $\frac{\frac{\sqrt{n}\left(Cn^{-2}+\frac{1}{n}\right)}{\sqrt{\frac{k}{n}\left(1-\frac{k}{n}\right)}}\frac{1}{\sqrt{2\pi}}\exp\left\{-\frac{\left[\frac{\sqrt{n}}{\sqrt{\frac{k}{n}\left(1-\frac{k}{n}\right)}}(\lambda-\pi)\right]^2}{2}\right\}}{\frac{1}{n}} \rightarrow_{a.s.} 0$. But rearranging this expression yields

$$\frac{\left(\frac{C}{n}+1\right)}{\sqrt{2\pi}}\frac{\sqrt{n}}{\sqrt{\frac{k}{n}\left(1-\frac{k}{n}\right)}}\exp\left\{-\frac{n}{2\frac{k}{n}\left(1-\frac{k}{n}\right)}(\lambda-\pi)^2\right\},$$

which converges (a.s.) to zero if and only if

$$\frac{\sqrt{n}}{\sqrt{\frac{k}{n}\left(1-\frac{k}{n}\right)}}\exp\left\{-\frac{n}{2\frac{k}{n}\left(1-\frac{k}{n}\right)}(\lambda-\pi)^2\right\} \rightarrow_{a.s.} 0. \tag{27}$$

Using the fact that the function above is continuous, $\sqrt{\frac{k}{n}\left(1-\frac{k}{n}\right)} \rightarrow_{as} \sqrt{\pi(1-\pi)}$, and

$$\frac{\sqrt{n}}{\sqrt{\pi(1-\pi)}}\exp\left\{-\frac{n}{2\pi(1-\pi)}(\lambda-\pi)^2\right\} \rightarrow_{a.s.} 0.$$

Thus, (27) is satisfied whenever $\lambda \neq \pi$. This establishes our desired result. ■

Proof of Theorem 6. Consider a history h^t with k high interpretations in n informative sub-histories. We want to show that there exists γ_H such that for $\frac{k}{n} > \gamma_H$ and n large, the individual will always play $\hat{\sigma} = H$. In order to obtain a contradiction, suppose this is not true and consider a deviation to $\hat{\sigma} = H$. The proof proceeds by obtaining an upper bound on the equilibrium payoffs, a lower bound on the payoff from deviating, and showing that the lower bound on the payoff from deviating is greater than the upper bound on the equilibrium payoffs. I will consider the case in which $\frac{k}{n} > \gamma_H$ (by symmetry, the other case is analogous).

The payoff from deviating is bounded below by the strategy in which the individual plays $\hat{\sigma} = H$ after every future low signal. In this case, all additional signals are uninformative and the individual gets an anticipatory payoff of $E\left[\theta - \frac{1}{2}|k+1,n\right]$ and an expected payoff from actions equal to $E\left[\theta - \frac{1}{2}|k,n\right]$. Thus, the payoff from deviating is bounded below by

$$E\left[\theta - \frac{1}{2}|k+1,n\right] + E\left[\theta - \frac{1}{2}|k,n\right] = \frac{2k+3}{n+2} - 1. \tag{28}$$

The payoff from playing any equilibrium strategy is bounded above by the payoff obtained if the individual could commit to interpreting every signal realistically. Since $\eta \to 0$, the actions taken under this interpretation strategy converge (almost surely) to $a = \begin{cases} 1 \text{ if } \theta > \frac{1}{2} \\ -1 \text{ if } \theta < \frac{1}{2} \end{cases}$, which yields an expected payoff of

$$2E\left[\left(\theta - \frac{1}{2}\right)\mathbf{1}\left(\theta \geq \frac{1}{2}\right) - \left(\theta - \frac{1}{2}\right)\mathbf{1}\left(\theta < \frac{1}{2}\right)|k,n\right],$$

where $\mathbf{1}(x)$ denotes the indicator function. Rearranging this expression, we obtain

$$\frac{2\left[\int_{\frac{1}{2}}^{1}\theta^{k+1}(1-\theta)^{n-k}d\theta - \int_{0}^{\frac{1}{2}}\theta^{k+1}(1-\theta)^{n-k}d\theta\right] + \int_{0}^{\frac{1}{2}}\theta^{k}(1-\theta)^{n-k}d\theta - \int_{\frac{1}{2}}^{1}\theta^{k}(1-\theta)^{n-k}d\theta}{\int_0^1\theta^k(1-\theta)^{n-k}d\theta}. \tag{29}$$

Next, we establish that, for $\frac{k}{n} \equiv \gamma \approx 1$ and n large enough, the expression in (28) is greater than the one in (29). By continuity, it suffices to check that (28) is strictly greater than the one in (29) for $\gamma = 1$ when n is large enough. Thus, it suffices to show that

$$\xi(n) \equiv \left(\frac{n+1}{n+2}\right)\int_0^1 \theta^n d\theta - 2\left(\int_{\frac{1}{2}}^1 \theta^{n+1} d\theta - \int_0^{\frac{1}{2}} \theta^{n+1} d\theta\right) - \int_0^{\frac{1}{2}} \theta^n d\theta + \int_{\frac{1}{2}}^1 \theta^n d\theta > 0,$$

for large n. It is straightforward to verify that $\xi(n) \neq 0$ and $\xi(n) \searrow 0$, which implies that there exists an $\bar{n}$ such that $n > \bar{n} \implies \xi(n) > 0$, concluding the proof. ■

Proof of Proposition 5. Let $h^{T-\tau}$ be a history with k high interpretations in n informative subhistories and suppose that $\frac{k}{n} \geq \frac{1}{2}$. Suppose the equilibrium assigns $\hat{\sigma}^*\left(h^{T-\tau}\right) = L$ and consider a deviation to $\hat{\sigma} = H$. A sufficient condition for this deviation to be profitable is that, for every continuation interpretation strategy after the deviation, the individual keeps playing action $a = 1$ and she weakly prefers to play this action (hence, this deviation has no costs in terms of suboptimal decision making). Since the individual observes τ additional signals, the individual keeps playing $a = 1$ in every continuation strategy and every continuation history after the deviation if she chooses $a = 1$ even after observing τ additional low signals and interprets them realistically:

$$\frac{k+1}{n+\tau} > \frac{1}{2}.$$

Moreover, she prefers action $a = 1$ for every continuation strategy and every continuation history if

$$\frac{k}{n+\tau} \geq \frac{1}{2}.$$

Combining these two conditions yields $\frac{k}{n+\tau} \geq \frac{1}{2}$.

Now let $\frac{k}{n} \leq \frac{1}{2}$, suppose the equilibrium assigns $\hat{\sigma}^*\left(h^{T-\tau}\right) = L$, and consider a deviation to $\hat{\sigma} = H$. Again, a sufficient condition for this deviation to be profitable is that, for every continuation interpretation strategy, the individual keeps playing $a = 0$ and she weakly prefers this action in every continuation history. She keeps playing $a = 0$ if she does so even after observing τ high signals and interpreting them realistically:

$$\frac{k+\tau+1}{n+\tau} < \frac{1}{2}.$$

Furthermore, she prefers action $a = 0$ in all continuation histories if

$$\frac{k+\tau}{n+\tau} \leq \frac{1}{2}.$$

Combining these two inequalities yields $\frac{k+\tau+1}{n+\tau} < \frac{1}{2}$. ■

Proof of Proposition 6. The proof follows the exact steps as the proof of Proposition 3. The only difference is on the bound of the discounted gain in self-image, which becomes

$$\alpha \sum_{t=0}^{T-1} \frac{\beta^t}{t+3} = \begin{cases} \alpha\left[\frac{1}{3} + \frac{\beta}{T+3}\left(\frac{1-\beta^T}{1-\beta}\right)\right] \text{ if } \beta < 1, \text{ and} \\ \alpha \sum_{t=3}^{T-1} \frac{1}{t} \text{ if } \beta = 1. \end{cases} \qquad ■$$

Proof of Proposition 7. The proof follows the exact steps as the proof of Proposition 4, except that the expected payoffs from self image become $\frac{\alpha}{3}\left(\frac{1-\beta^T}{1-\beta}\right)$ if the individual interprets a low signal realistically at state $(0,0)$, and $\frac{2\alpha}{3}\left(\frac{1-\beta^T}{1-\beta}\right) - \frac{\kappa}{2}$ if she deviates and chooses $\hat{\sigma} = H$. Deviating is not profitable since $\kappa \geq \frac{2\alpha}{3}\left(\frac{1-\beta^T}{1-\beta}\right) \implies \frac{\alpha}{3}\left(\frac{1-\beta^T}{1-\beta}\right) \geq \frac{2\alpha}{3}\left(\frac{1-\beta^T}{1-\beta}\right) - \frac{\kappa}{2}$ ■

References

[1] Acemoglu, D., V. Chernozhukov, and M. Yildiz (2009), “Fragility of Asymptotic Agreement under Bayesian Learning,” Mimeo., MIT.

[2] Akerlof, G. A. and W. T. Dickens (1982). “The Economic Consequences of Cognitive Dissonance,” American Economic Review, 72(3), 307-319.

[3] Ali, N. (2009), “Learning Self-Control,” Mimeo., UCSD.

[4] Alloy and Ahrens (1987). “Depression and Pessimism for the Future: Biased Use of Statistically Relevant Information in Predictions for Self Versus Others,” Journal of Personality and Social Psychology.

[5] Bahrick, H. P., L. K. Hall, and S. A. Berger (1996). “Accuracy and distortion in memory for high school grades,” Psychological Science, 7, 265-271.

[6] Barberis, N., A. Shleifer, and R. Vishny (1998). “A model of investor sentiment,” Journal of Financial Economics, 49 (3), 307-343.

[7] Beck, A.T. (1967). Depression: Clinical, experimental, and theoretical aspects, Harper & Row, New York.

[8] Benabou, R. (2008a). “Groupthink: Collective Delusions in Organizations and Markets,” Mimeo., Princeton University.

[9] Benabou, R. (2008b). “Ideology,” Journal of the European Economic Association, 6, 321–352.

[10] Benabou, R. and G. Laroque (1992). “Using Privileged Information to Manipulate Markets: Insiders, Gurus, and Credibility,” Quarterly Journal of Economics, 107, 921-948.

[11] Benabou, R. and J. Tirole (2002). “Self-Confidence and Personal Motivation,” Quarterly Journal of Economics, 117(3), 871-915.

[12] Benabou, R. and J. Tirole (2004). “Willpower and Personal Rules,” Journal of Political Economy, 112 (4), 848-886.

[13] Benabou, R. and J. Tirole (2006a). “Belief in a Just World and Redistributive Politics,” Quarterly Journal of Economics, 121(2), 699-746.

[14] Benabou, R. and J. Tirole (2006b). “Incentives and Prosocial Behavior,” American Economic Review, 96(5), 1652-1678.

[15] Benabou, R. and J. Tirole (2006c). “Identity, Dignity and Taboos: Beliefs as Assets,” Mimeo., Princeton University and Université de Toulouse.

[16] Bernheim, B. D. and R. Thomadsen (2005). “Memory and Anticipation,” Economic Journal, 115, 271–304.

[17] Bhaskar, V., G. Mailath, and S. Morris (2009). “A Foundation for Markov Equilibria in Infinite Horizon Perfect Information Games,” Mimeo., UCL, University of Pennsylvania, and Princeton University.

[18] Bodner, R. and D. Prelec. “Self-signaling in a neo-Calvinist model of everyday decision making,” in Psychology and Economics, Vol. II. Brocas and J. Carillo (eds.), Oxford University Press, 2002.

[19] Brocas, I. and J. Carrillo, (2009). “"From Perception to Action: an Economic Model of Brain Processes,” Mimeo., USC.

[20] Carrillo, J., and T. Mariotti (2000). “Strategic Ignorance as a Self-Disciplining Device,” Review of Economic Studies, 66 , 529–544.

[21] Di Tella, R., S. F. Galiani, and E. S. Schargrodsky (2007). "The Formation of Beliefs: Evidence from the Allocation of Land Titles to Squatters," Quarterly Journal of Economics, 122(1).

[22] Edwards, W. (1968). "Conservatism in human information processing," In: B. Kleinmutz (Ed.), Formal Representation of Human Judgment. (pp. 17-52). New York: John Wiley and Sons.

[23] Evans, J. (1989). Bias in human reasoning: Causes and consequences. Hillsdale, NJ: Erlbaum.

[24] Fudenberg, D. and J. Tirole (1991). Game Theory. Cambridge, MA: MIT Press.

[25] Gennaioli, N. and A. Shleifer (2010). "What Comes to Mind," Quarterly Journal of Economics, forthcoming.

[26] Gollwitzer, P. M., Earle, W. B., and Stephan, W. G. (1982). "Affect as a determinant of egotism: Residual excitation and performance attributions," Journal of Personality and Social Psychology, 43, 702-709.

[27] Gottlieb, D. (2009). "Imperfect Memory and Choice under Risk," Mimeo. Princeton University and Wharton School of Business.

[28] Ghosh, J.K. and R.V. Ramamoorthi (2003). Bayesian nonparametrics, New York: Springer-Verlag.

[29] Greene, E. (1981). "Whodunit? Memory for evidence in text," American Journal of Psychology, 94, 479-496.

[30] Johnson, R.A. (1970) "Asymptotic Expansions Associated with Posterior Distributions," The Annals of Mathematical Statistics, 41(3), 851-864.

[31] Kopczuk, W. and J. Slemrod (2005). "Denial of Death and Economic Behavior," Advances in Theoretical Economics, 5(1), Article 5.

[32] Kuiper, N. A. (1978). Depression and causal attributions for success and failure. Journal of Personality and Social Psychology, 36, 236–246.

[33] Kunda, Z. (1990). "The case for motivated reasoning," Psychological Bulletin, 108, 480-498.

[34] Kuran, T. (1993). "The Unthinkable and the Unthought," Rationality and Society, 5, 473-505.

[35] Lord, C. G., L. Ross, and M. R. Lepper (1979). "Biased assimilation and attitude polarization: The effects of prior theories on subsequently considered evidence," Journal of Personality and Social Psychology, 37 (11). 2098–2109.

[36] MacDonald, T. K., and M. Ross (1999). "Assessing the accuracy of predictions about dating relationships: How and why do lovers' predictions differ from those made by observers?," Personality and Social Psychology Bulletin, 25, 1417-1429.

[37] Madarász, K. (2009). "Information Projection: Model and Applications ," Mimeo., LSE.

[38] Maskin, E. and J. Tirole (2001). "Markov perfect equilibrium: I. Observable Actions," Journal of Economic Theory, 100(2), 191-219.

[39] Mather, M., E. Shafir, and M.K. Johnson (2003). "Memory and Remembering chosen and assigned options," Memory & Cognition, 31 (3), 422-433.

[40] Montier, J. (2007). "The Limits to Learning," In: Behavioural Investing: A Practitioner's Guide to Applying Behavioural Finance, J. Montier, Wiley Finance.

[41] Morgan, J. and P. C. Stocken (2003). "An analysis of stock recommendations," RAND Journal of Economics, 34(1), pp. 183-203.

[42] Moore, D. A. (2007). “Not so above average after all: When people believe they are worse than average and its implications for theories of bias in social comparison,” Organizational Behavior and Human Decision Processes, 102(1), 42-58.

[43] Morris, S. (2001). “Political Correctness,” Journal of Political Economy 109, 231-265.

[44] Nickerson, R. S. (1998). “Confirmation bias: A ubiquitous phenomenon in many guises,” Review of General Psychology, 2(2), 175-220.

[45] Oswald, M. E. and S. Grosjean (2004). “Confirmation Bias,” In: R. F. Pohl. Cognitive Illusions: A Handbook on Fallacies and Biases in Thinking, Judgement and Memory. Hove, UK: Psychology Press.

[46] Ottaviani, M. and P. N. Sørensen (2006a). “Reputational Cheap Talk,” RAND Journal of Economics, 37, No. 1, pp. 155-175.

[47] Ottaviani, M. and P. N. Sørensen (2006b). “Professional Advice,” Journal of Economic Theory, 126(1), pp. 120-142.

[48] Peterson, C., Schwartz, S. M., & Seligman, M. E. p. (1981). Self-blame and depressive symptoms. Journal of Personality and Social Psychology, 41, 253–259.

[49] Piccione, M. and Rubinstein, A. (1997). “On the interpretation of decision problems with imperfect recall,” Games and Economic Behavior, 20, 3–24.

[50] Rabin, M. (1994). “Cognitive Dissonance and Social Change,” Journal of Economic Behavior and Organization, 23, 177-194.

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

[52] Sanbonmatsu, D.M., S. S. Shavitt, S.J. Sherman, and D.R. Roskos-Ewoldsen (1987). “Illusory correlation in the perception of performance by self or a salient other,” Journal of Experimental Social Psychology, 23 (6), 518–543.

[53] Schelling, T. (1985). “The Mind as a Consuming Organ,” In: J. Elster, ed. The Multiple Self. New York: Cambridge University Press.

[54] Schwartzstein, J. (2010). “Selective Attention and Learning,” Mimeo., Harvard University.

[55] 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.

[56] Seligman, Martin E.; Abramson, Lyn Y.; Semmel, Amy; von Baeyer, Carl (1979) Depressive attributional style. Journal of Abnormal Psychology. Vol 88(3), 242-247.

[57] Shepperd, J., W. Malone, and K. Sweeny, (2008). “Exploring Causes of the Self-serving Bias,” Social Personality Psychology Compass, 2, 895-908.

[58] Smith, J. (2009), “Imperfect Memory and the Preference for Increasing Payments,” Journal of Institutional and Theoretical Economics, 2009, 165(4): 684-700.

[59] Smith, J. (2009), “Cognitive Dissonance and the Overtaking Anomaly: Psychology in the Principal-Agent Relationship,” Journal of Socio-Economics, 2009, 38(4): 684-690.

[60] Sobel, J. (1985), “A Theory of Credibility,” Review of Economic Studies, 52, 557-573.

[61] Taylor, S. E. and P. M. Gollwitzer (1995). “Effects of Mindset on Positive Illusions,” Journal of Personality and Social Psychology, 69(2), 213-226.

[62] Tennen and Herzberger (1987). “Depression, Self-Esteem, and the Absence of Self-Protective Attributional Biases,” Journal of Personality and Social Psychology.

[63] Wilson, A. (2005), "Bounded Memory and Biases in Information Processing," Mimeo. New York University.

[64] Zuckerman, M. (1979). "Attibution of success and failure revisited, or: The motivational bias is alive and well in attribution theory," Journal of Personality, 47, 245-287.