

I'll See It When I Believe It — A Simple Model of Cognitive Consistency

By

Leeat Yariv

February 2002

COWLES FOUNDATION DISCUSSION PAPER NO. 1352



COWLES FOUNDATION FOR RESEARCH IN ECONOMICS

YALE UNIVERSITY

Box 208281

New Haven, Connecticut 06520-8281

<http://cowles.econ.yale.edu/>

I'LL SEE IT WHEN I BELIEVE IT - A SIMPLE MODEL OF COGNITIVE CONSISTENCY*

Leeat Yariv[†]

First Version: June 2000, This Version: January 2002

Abstract

Psychological experiments demonstrate that people exhibit a taste for consistency. Individuals are inclined to interpret new evidence in ways that confirm their pre-existing beliefs. They also tend to change their beliefs to enhance the desirability of their past actions. I present a model that incorporates these effects into an agent's utility function and allows me to characterize when: (i) agents become under- and over-confident, (ii) agents exhibit excess stickiness in action choices, (iii) agents prefer less accurate signals, and (iv) agents are willing to pay in order to forgo signals altogether. Applications to political campaigns and investment decisions are explored.

Key words: Belief utility, cognitive dissonance, confirmatory bias, overconfidence, selective exposure.

Journal of Economics Literature Classification Numbers: C90, D72, D83, D91, M30.

I am extremely grateful to Muriel Niederle for the many coffees and chats that inspired this project and later helped keep it going. I would like to express my gratitude to Drew Fudenberg, David Laibson, and Eric Maskin for endless conversations, helpful comments and suggestions, and support. I've also benefitted from talking with Alberto Alesina, Ed Glaeser, Orit Kedar, Ted Miguel, Paul Milgrom, Markus Mobius, Al Roth, Michael Schwartz, and Steve Tadelis. I am very thankful to the Alfred P. Sloan foundation and Eliot dissertation completion committee for generous financial support.

Department of Economics, Yale University and Department of Economics, University of California at Los Angeles.
E-mail: leeat.yariv@yale.edu

“I’ll see it when I believe it.”

-slip of tongue by experimental social psychologist Thane Pittman.

“Nothing is quite as difficult as not deceiving oneself.”

-Ludwig Wittgenstein, 1889-1951.

I. Introduction

Consider someone who buys an expensive car and later discovers that it is uncomfortable on long drives. Disagreement, or *dissonance*, exists between their beliefs that they have bought a good car and that a good car should be comfortable. Dissonance could be reduced by deciding this discomfort does not matter since the car is mainly used for short trips (decreasing the importance of the dissonant belief) or focusing on the car’s strengths such as safety, appearance, and handling (thereby adding more consonant beliefs). The dissonance could also be eliminated by getting rid of the car, but this behavior is a lot harder than changing beliefs. Moreover, if ambiguous information about the car arrives (a friend comments “this looks like a really interesting car”), it is more likely to be interpreted as supporting the belief that the car is good (“the friend means the car looks really cool and unique”) rather than countering it (“what is an interesting looking car? The friend must mean it looks like a real lemon. Maybe the fact it’s uncomfortable on long drives shows.”).

This example illustrates two cognitive biases that are consistently observed in agents making decisions under uncertainty: cognitive dissonance and confirmatory bias. *Cognitive dissonance* states that after having taken an action people tend to change their beliefs about the relative agreeableness of this action. *Confirmatory bias* refers to the phenomenon of people interpreting new evidence in ways that confirm their current beliefs.

The abundant psychological literature on these two biases (overviewed in Section II) indicates that people behave as if they have a taste for: 1. consistency between the action taken and the belief held at each point in the decision process and 2. consistency between beliefs in different stages of the decision process.

This paper proposes a model in which utility functions depend directly on beliefs. The suggested framework helps explain an assortment of experimental observations related to information economics. In particular, I characterize when agents become overconfident or underconfident, and identify cases in which agents are likely to persist with the action they chose initially. Furthermore, I show that agents may sometimes prefer less accurate signals over more accurate ones and less signals to more signals.

I consider a separable utility function that is comprised of two terms. The first is *instrumental utility*, which coincides with the utility considered in standard economic analysis. The second term

is *belief utility*, corresponding to the direct utility from the sequence of beliefs adopted by the agent. This form of utility functions captures the tradeoff between making optimal choices according to standard analysis (maximizing the instrumental utility) and being consistent (maximizing the belief utility). In addition, the functional form used throughout the paper belongs to a class of utility functions that have an axiomatic foundations grounded in some of the psychological observations surveyed in Section II (see Yariv [2001]).

I look at a dynamic process in which an agent tries to guess the state of nature out of two possible states. At each stage the agent receives a (probabilistic) signal indicating what the state of the world is. The agent has to then choose a belief and an action.

As compared with standard models of decisions under uncertainty, there are two novel components in this setup: the direct dependency of utility on beliefs and the ability of the agent to *choose* her beliefs. In other words, the current model treats beliefs as actions.¹

Once agents choose their beliefs, their level of sophistication, or introspection, as measured by the extent to which they perceive correctly their future behavior, plays a crucial part in the outcomes of the decision process. I parametrize introspection with one (continuous) variable that denotes the weight the agent thinks she will give the belief utility in the future.²

When agents have a taste for consistency and are sufficiently forward looking (introspective), accuracies of future signals may affect their current choices of beliefs. For example, if a future signal is very accurate, agents might prefer to choose a more conservative belief in the present that would allow them to follow the accurate future signal without losing much belief utility. Proposition 1 specifies the circumstances under which an agent will choose a belief too extreme or too conservative with respect to the Bayesian belief. Simply put, if agents put sufficient weight on being consistent, and have a sufficiently accurate initial signal, then they are overconfident when the initial signal supports their prior belief and underconfident when their initial signal opposes their prior belief.

Biases in beliefs have a direct effect on the actions taken. Intuitively, consider an agent who does not fully perceive her future taste for consistency (i.e., has limited introspection). She might choose extreme beliefs at the outset in order to get immediate belief utility, erroneously predicting that she will not distort much her future Bayesian posterior. However, once a later period arrives, the agent does care about consistency, does not update her beliefs sufficiently, and as a consequence

¹Note that everyday language suggests the resemblance of beliefs and standard actions. E.g., we talk about “holding beliefs,” “giving up beliefs,” “changing beliefs,” etc.

²For example, *forward looking* agents are ones in which the weight the agent thinks she’ll give the belief utility is the actual true weight. Such agents are hyper-rational in that they fully perceive, and take into account, the effects of current choices of actions and beliefs on all future ones. On the other end of the spectrum, *naive optimists* consider the effects of their current choices on future ones, but assume they will not have any non-instrumental motivations in the future. They are time inconsistent in that they mistakenly believe that they will behave as perfect Bayesians in the future.

persists choosing her past action. In such situations, history fully determines the path of actions to be taken, even when pure statistical inference would not provide a justification for this stickiness in actions. Proposition 2 gives a formal characterization of situations in which such excess persistence arises.

A Bayesian econometrician trying to fit a utility function to the agent's choices, observing the available signals, would thus see a seemingly increasing taste for the initial actions taken. Differently put, excess persistence in action choices is observationally equivalent to habit formation. This equivalence is formalized in Section IV.

Looking at expected utility levels, I derive general results taking two points of view. The first is paternalistic. If an outsider can fully predict the agent's behavior, what kind of signals would she choose in order to maximize the agent's current expected utility. The other point of view is non-paternalistic. I ask what kind of signals the agent herself would prefer, taking into account her possible limited introspection. In both cases it turns out that expected utility may be lower when signals are more accurate (Propositions 3 and 4) or more abundant (Proposition 5). Intuitively, a more accurate signal may raise instantaneous utility, but may at the same time cause an agent to choose more extreme beliefs, hence hindering the possibility of assimilating other pieces of information without forgoing a greater cost in terms of the belief utility. When the latter effect dominates, the agent may prefer less accurate signals, or less signals altogether.

Overconfidence, underconfidence, excess persistence in action choices, and habit formation are phenomena that have been identified in economic behavior (see Section II). The taste for less accurate information, as captured in Propositions 4 and 5, correspond to observations that have not yet been tested for in the field or, experimentally, in the lab. Hence, the current framework both organizes a few observed phenomena under the same theoretical umbrella, as well as provides new predictions that can serve as a natural testing ground.

The model has a wide range of applications. In political economy, the notion of selective exposure is folk wisdom (e.g., political voters expose themselves to propaganda favoring their side - see Sears and Freedman [1967]). Campaigns and political debates may be viewed as signals which the voters can choose to observe or ignore. In marketing, the results shed light on ways to exploit agents who care about consistency (an example is the foot-in-the-door technique which gets consumers to make a big commitment following a small one - see, e.g., Freedman and Fraser [1966]). In business management, the model partially explains the tendency of managers to persist with unsuccessful policies (see Staw [1976]). In finance, the model predicts a bias towards early investments and continuous use of sub-optimal sources of information, as well as under and overconfidence concerning investment decisions (for an empirical overview, see Shleifer [2000]). In labor, the model predicts overconfidence of workers in hazardous occupations (see Akerlof and Dickens [1982] and references

therein). Some of these applications are explored in more detail in Section V.

The structure of the paper is as follows. Section II provides some psychological evidence for confirmatory bias and cognitive dissonance, as well as reviews the existing literature in economics connecting to the current paper. Section III contains the general setup of the dynamic decision problem treated. Section IV presents the main results for the restricted setup. Sections IV.A and IV.B analyze the choice of beliefs and actions, while Section IV.C characterizes the dependency of expected utilities on the information structure. Section V discusses applications of the model. Section VI contains some results on strategic interactions between agents with preferences for consistency. Section VII concludes. Most of the proofs are relegated to the appendices. Appendix A contains the proofs for the results in the applied part of the paper. Appendix B contains some special results for forward looking agents.

II. Psychological Evidence and Related Literature

Confirmatory bias is identified with the phenomenon of people interpreting new evidence in ways that confirm their current beliefs. A large body of literature illustrates this phenomenon. In an important early paper, Oskamp [1965] gave practicing and training clinical psychologists a gradual description of a case. Each therapist was presented with four sets of cumulatively increasing amounts of information as the basis for making his or her decisions. Oskamp discovered that beyond some early point in the information-gathering process, predictive accuracy reaches a ceiling. Nevertheless, confidence in one's decisions continues to climb steadily as more information is obtained. In a similar spirit, Darley and Gross [1983] demonstrated that teachers misread performance of pupils as supporting their initial impressions of those pupils.

Frank and Gilovich [1988] first asked a group of respondents to rate the appearance of professional football and hockey uniforms. Subjects judged those uniforms that were at least half black to be the most “bad,” “mean,” and “aggressive” looking. These perceptions influenced, in turn, how specific actions performed by black uniformed teams were viewed. The authors showed groups of trained referees one of two videotapes of the same aggressive play in a football scrimmage, one with the aggressive team wearing white and one with it wearing black. The referees who saw the black-uniformed version rated the play as much more aggressive and more deserving of a penalty than those who saw the white-uniformed version. The referees “saw” what this common negative association led them to expect to see. As a result of this bias, it is not surprising to learn that teams that wear black uniforms in these two sports have been penalized significantly more than average during the last two decades.

There is some evidence indicating that the confirmatory bias phenomenon is hard-wired in

our brains. The McGork Effect in linguistics refers to the observation that when subjects are shown a video of a person saying d sounds (lips separated in pronunciation), and the accompanying soundtrack is that of b sounds (lips closed in pronunciation), subjects report hearing g (as in “great”) sounds. This effect vanishes immediately when subjects are asked to shut their eyes. Thus, it seems as if the brain goes automatically into a process of resolving an inconsistency between visual and aural stimuli (see Fromkin [2000] for a more elaborate description of the experiment).

Cognitive Dissonance is a broad bias, which asserts that after having chosen an action people tend to change their beliefs about the relative agreeableness of this action. Festinger [1957] was one of the first to regard attitudes held by a single individual to exist in a state of tension, which he termed *dissonance*. This occurs when an individual does something that follows neither from the attitudes the person holds nor from some extrinsic force such as the expectation of reward. Festinger showed that in such a situation people can be expected to move their beliefs into line with their behavior. As an illustration, if an individual is maneuvered into delivering a speech that happens not to reflect the person’s prior beliefs, and if the person is paid little or nothing for doing so, the person’s expressed attitudes move in the direction of the position taken in the speech. This movement is blocked if the person is paid a substantial amount for delivering the speech. In the latter case, giving the speech is highly consistent with the payment and the person recognizes the lack of relation between prior beliefs and what was said. Similarly, students who are paid very little to perform a boring task describe it to their classmates as relatively interesting. A few experimenters used a similar framework to illustrate justification of cruel behavior. For example, Glass [1964] reported that students who were asked to give electrical shocks to victims subsequently lowered their opinions of their victims (a good summary of seminal experiments on cognitive dissonance appears in Aronson [1969] and Nisbett and Ross [1991]).

In addition, experiments show that people with the same information differ systematically in their beliefs depending on their theories, following an action. For instance, Brehm [1956] asked women to rate the worthiness of two appliances. The subjects then chose one of the appliances and were given their choice wrapped. A few minutes later, and before the boxes were unwrapped, the women were given a second evaluation of the two appliances. These evaluations changed systematically in favor of the appliance that had been chosen. Similarly, Knox and Inkster [1968] interviewed bettors at a race track and showed that people just leaving the betting window placed much higher odds on “their horse” than people in the queue who hadn’t bought a betting ticket yet.

There have been several attempts in the literature to incorporate beliefs directly into the agent’s utility function. Bodner and Prelec [1997] and Akerlof and Kranton [2000] introduce the ideas of self-signaling and identity. These authors consider agents who value their beliefs about themselves

and thus choose actions that not only maximize some instrumental utility, but also a utility that reflects their self-perceptions. However, these models consider agents who choose only actions (and not beliefs) and resolve uncertainties (about their own types) by standard Bayesian updating, using their own actions as signals. The present model goes a step further, by allowing for external signals about uncertainty that pertain not only to the individual herself. Furthermore, I model how an agent will endogenously misread signals and choose her beliefs.

More recently, Yariv [2001] provided an axiomatic foundation for considering generalized discounted utility functional forms that contain beliefs as arguments. The functional form used in the current paper is, in fact, a special case of that characterized by Yariv.

On the more applied aspect, Akerlof and Dickens [1982] proposed a simple model of cognitive dissonance relevant to workers in hazardous occupations. They looked at a two period model in which cost-effective safety equipment is available only in the second period and showed that cognitive dissonance implies that workers will convince themselves in the first period that they are not unsafe without the safety equipment. Formally, workers change their probabilistic assessments of the hazards of not being protected by an exogenously given amount. As a consequence, agents will not purchase the equipment in the second period.

Rabin and Schrag [1999] present a model of confirmatory bias in which, with some exogenously given probability, people misread signals that contradict their current beliefs. Some of my results concerning overconfidence resemble Rabin and Schrag's in spirit. However, the model provided in this paper differs from theirs, as well as from Akerlof and Dickens', in that I endogenize the probability with which an agent misinterprets information available to her, and the extent to which she does so. Moreover, I make predictions concerning the quality of signals agents will choose to obtain.

Benabou and Tirole [2001] and Koszegi [1999] propose models that explain overconfidence. They assume very specific functional forms and a dynamics in which only one action is taken. Furthermore, the agents' beliefs are a choice variable in the sense that the agents choose when to stop receiving information. Unlike the current setting, once the stopping time is set, Bayesian updating is used to construct beliefs throughout the decision process.

To summarize, the link between all of these papers is the acknowledgement of motivated distortionary beliefs. The model put forth in this paper is different in that it admits the endogenous choice of beliefs - given an information structure, agents choose both actions and beliefs.

III. Setup

A. Underlying Framework

I consider a world with two possible states $\theta \in \Omega = \{L, R\}$, with a-priori probability $\Pr(\theta = L) = p > \frac{1}{2}$. That is, (without loss of generality) I assume agents believe L is more likely to be the state of the world.

Beliefs are then probability assessments of the state of the world θ being L .

I assume the set of actions coincides with the states of nature. That is, $A = \Omega = \{L, R\}$.

At each stage, the agent tries to guess the state of the world. She receives (instrumental) utility if her guess is correct. Formally, the agent can choose an action $a \in \{L, R\}$, which delivers stage utilities of $u(a = \theta) = 1$ and $u(a \neq \theta) = 0$.

The states of the world can be viewed as a metaphor for an assortment of economic and non-economic variables. For example, L and R can stand for the superiority of different political candidates (*Left* or *Right*), fit for different professions (weight *Lifting* or *Research*), personal matches (*Leo* or *Rob*), different business or financial policies (such as investing or not investing), etc.

When choosing an action, I assume that agents have to act in a way that is consistent with their beliefs. Hence, agents' decisions satisfy:

$$a_t \equiv a(\mu_t) = \begin{cases} L & \mu_t > \frac{1}{2} \\ L \text{ or } R & \mu_t = \frac{1}{2} \\ R & \mu_t < \frac{1}{2} \end{cases} .$$

Using the above metaphors, a voter who believes that *Left* is more likely to provide better leadership will vote *Left*, an undergraduate who believes weight *Lifting* is likely to be a better fit for her than *Research* will choose an athletic career over an academic one, etc.

The utility of the agent is comprised of two parts: The first is instrumental utility, which concerns the per period payoffs of actions. The second is a belief utility, which corresponds to the utility an agent gets from changes in beliefs. That is, total utility at time t for a T period process is of the form:

$$U_t = \sum_{\tau=t}^T \delta^{\tau-t} [u(a_\tau) + \gamma v(\mu_{\tau-1}, a_\tau, \mu_\tau)]$$

where δ is the discount factor (which will usually be assumed to equal 0 or 1), u denotes the instantaneous instrumental utility, and v denotes the instantaneous belief utility.

An axiomatic foundation for generalized discounted utility functional forms which have both

actions and beliefs as arguments is provided in Yariv [2001].

Since the action is assumed to be completely determined by the belief at each stage, i.e., $a_t = a(\mu_t)$, v is in fact a function of only two arguments and one can write:

$$U_\tau = \sum_{\tau=t}^T \delta^{\tau-t} [u(a_\tau) + \gamma v(\mu_{\tau-1}, \mu_\tau)].$$

At each stage $t \geq 1$, the agent receives a signal s_t of accuracy q_t . That is, $\Pr(s_t = \theta) = q_t$. These signals are conditionally independent across time.

At each stage the agent chooses an action and belief so as to maximize her expected total utility. Since I am assuming that actions need to be consistent with the beliefs, they are fully determined by the beliefs. In particular, the agent has essentially one choice variable which is the beliefs. It is important to note here that for a finite set of actions and a continuum of beliefs, specifying the action doesn't indicate which belief was chosen. That is, the model cannot be translated into a standard model in which the actions are chosen and the beliefs are implied by the actions as in the standard framework.³

Remark 1: Once the agent chooses her beliefs, the assumption that actions are consistent with the beliefs is quite strong. This can be an outcome of a utility specifying a very high cost of choosing an action that doesn't correspond to the belief held (for example, when the agent needs to justify her actions - see Tetlock et al. [1989]).

If the belief held at time t is μ_t and signal s_{t+1} of accuracy q_{t+1} is observed, then the updated posterior, according to Bayes' rule, is given by:

$$\mu_{t+1}^B = \begin{cases} \frac{\mu_t q_t}{\mu_t q_t + (1-\mu_t)(1-q_t)} & s_t = L \\ \frac{\mu_t(1-q_t)}{\mu_t(1-q_t) + (1-\mu_t)q_t} & s_t = R \end{cases}.$$

Remark 2: In the first stage, Bayesian updating yields:

$$\mu_1^B = \begin{cases} \frac{pq_1}{pq_1 + (1-p)(1-q_1)} & s_1 = L \\ \frac{p(1-q_1)}{p(1-q_1) + (1-p)q_1} & s_1 = R \end{cases}$$

³We could think of the agent as picking *accuracies* of signals and using only Bayes' rule to update. The model would then predict that an agent would not always pick the most accurate signal. Such a model is still slightly more limiting than ours since it doesn't allow for the *asymmetric* tilting of beliefs (towards those of the previous period). Put another way, while an agent might be able to choose between coins of different accuracies, she would not be able to choose their realizations.

whereas in the second stage,

$$\mu_2^B(\mu_1) = \begin{cases} \frac{\mu_1 q_2}{\mu_1 q_2 + (1-\mu_1)(1-q_2)} & s_2 = L \\ \frac{\mu_1(1-q_2)}{\mu_1(1-q_2) + (1-\mu_1)q_2} & s_2 = R \end{cases}$$

Note that while μ_1^B is the true Bayesian posterior after receiving s_1 , $\mu_2^B = \mu_2^B(\mu_1)$ is not necessarily the posterior a Bayesian agent would form after observing signals s_1 and s_2 . It is the Bayesian posterior at stage 2 when the prior is the adopted μ_1 and the signal s_2 is observed. μ_2^B is the fully Bayesian posterior only when the agent doesn't modify her first stage beliefs, i.e., when $\mu_1 = \mu_1^B$ and $\mu_2^B = \mu_2^B(\mu_1^B)$.

At each stage t , agents remember only their previous action a_t and belief μ_t .

I could assume that agents remember a longer block of actions and beliefs. This would not change the qualitative results, but would make the analysis more complicated algebraically.

The assumption that agents remember previous beliefs rather than previous signals and corresponding accuracies is grounded in some experimental observations. Indeed, there is evidence in the physiological psychology literature that people remember the theories they constructed better than the hard evidence upon which these beliefs were created (for a good overview of the field, see Schacter [1996]).

I assume agents use μ_t^B when choosing their beliefs and actions. I.e., agents calculate the posterior accurately and then choose, simultaneously, a belief and a corresponding action that balance the trade-off between the instrumental utility and the belief utility. This assumption is made to minimize the distance between the current model and the standard one. Once preferences are specified, the agent maximizes her well-being using all the information and statistical tools she has at hand, as is commonly assumed. Nonetheless, I view this as a strong assumption that serves as a benchmark, not necessarily as an accurate depiction of the cognitive processes that go on.⁴

In this paper I analyze two stage processes. This allows me to determine the effects of signals' accuracies on the decision outcomes, with the least presentational complexity.⁵

⁴In this framework I assume the agent solves at each stage t a problem of the sort $\max_{\mu_t, a_t} E_{\mu_t^B} U_t$. The other extreme is assuming the agent solves at time t the problem $\max_{\mu_t, a_t} E_{\mu_t} U_t$, in which case external signals have no informative value at all. One could look at parametric cases in between these two extremes. I.e., agents who solve problems of the form $\max_{\mu_t, a_t} E_{\alpha \mu_t^B + (1-\alpha) \mu_t} U_t$ where $0 \leq \alpha \leq 1$. This type of analysis is left for future research.

⁵Looking at only two stages turns out to be problematic since there are end effects that are confounded with the effects driven by the taste for consistency in this model. Looking at longer processes is possible, leads to similar qualitative results, and is far more cumbersome.

B. Summary of the Environment

There are two novel components in this setup relative to standard models of decision under uncertainty. First, the utility depends directly on beliefs. Second, the agent *chooses* her beliefs. Put differently, the current model treats beliefs as actions.

The timeline of the process is as follows:

- $t = 0$: prior belief $\mu_0 = p$, action a_0 chosen
- $t = 1$: signal s_1 given, belief μ_1 and corresponding action a_1 chosen
- $t = 2$: signal s_2 given, belief μ_2 and corresponding action a_2 chosen

The accuracies of the signals are given by:

$$\Pr(s_1 = \theta) = q_1 \quad \text{and} \quad \Pr(s_2 = \theta) = q_2$$

I assume the prior $p > \frac{1}{2}$ so that $a_0 = L$.

At each stage the agent bases her decisions on her past belief and current signal realization.

C. Agent Types

As described above, the agent's objective function at each stage t is:

$$\sum_{\tau=t}^T \delta^{\tau-t} [u(a(\mu_\tau)) + \gamma v(\mu_{\tau-1}, \mu_\tau)]$$

where $0 \leq \delta, \gamma \leq 1$.

In later parts of the paper agents' ability to predict their future preferences, namely predict the extent to which they will have a desire to behave consistently rather than maximize solely the instrumental utility, will play an important role in their choice of information structures. Throughout the paper, I will often term this ability to predict future preferences introspection.

In order to capture different levels of agents' introspection, I consider a scenario in which the agent has an objective function as above, but believes that for all stages $\tilde{t} > t$ the maximization problem is:

$$\max_{\mu_{\tilde{t}}} E_{\mu_{\tilde{t}}} \sum_{\tau=\tilde{t}}^T \delta^{\tau-t} [u(a(\mu_\tau)) + \tilde{\gamma} v(\mu_{\tau-1}, \mu_\tau)]$$

where $0 \leq \tilde{\gamma} \leq \gamma$.

That is, at each period, the agent thinks that from next period and on the relative weight of

the belief utility v will be only $\tilde{\gamma}$, and that she will care more about the instrumental utility than she does in the current period.

There are three special (extreme) cases of this formalism that are worth noting:

1. *Myopic agents* correspond to $\delta = 0$ (in which case the value of $\tilde{\gamma}$ is irrelevant to the agent's behavior). These agents form their beliefs so as to maximize their one period utility. They are not aware of the impact of their current bias on their future utility, and they do not take into account their future instantaneous utilities.

At each stage t , a myopic agent solves the following maximization problem:

$$\max_{\mu_t} E_{\mu_t^B} [u(a(\mu_t)) + \gamma v(\mu_{t-1}, \mu_t)]$$

2. *Forward looking agents* correspond to $\gamma = \tilde{\gamma}$. These agents take the whole future maximization problem into account, and do so correctly. These are the most sophisticated agents that we consider. They understand their own biases and consider them in their prediction of expected future utility.⁶

A forward looking agent solves at each stage t the following maximization problem:

$$\max_{\mu_t} E_{\mu_t^B} \sum_{\tau=t}^T \delta^{\tau-t} [u(a(\mu_\tau)) + \gamma v(\mu_{\tau-1}, \mu_\tau)]$$

$$\text{s.t. } \mu_{\tau'}(s_{\tau'}) \in \arg \max_{\mu} E_{\mu_{\tau'}^B(s_{\tau'})} \sum_{\tau=\tau'}^T \delta^{\tau-\tau'} [u(a(\mu_\tau)) + \gamma v(\mu_{\tau-1}, \mu_\tau)] \quad \forall \tau' > t, \forall s_{\tau'}.$$

3. *Naive optimists* correspond to $\tilde{\gamma} = 0$ - they take into account the effect of current belief choices on future choices of actions, but assume that all future beliefs will be chosen according to Bayes' rule. These agents are in-between the myopic and the forward looking in their levels of sophistication. They do consider the future, but do so inaccurately.

The maximization problem the naive optimist solves at each stage t is given by:

$$\max_{\mu_t} E_{\mu_t^B} \sum_{\tau=t}^T \delta^{\tau-t} [u(a(\mu_\tau)) + \gamma v(\mu_{\tau-1}, \mu_\tau)]$$

⁶A forward looking agent is sophisticated in the sense that she understands her own tendencies to bias her beliefs towards consistency. Gilovich [1991] presents an interesting quote by Charles Darwin indicating at least his awareness to his biases in that realm: "...[I] followed a golden rule, namely that whenever a new observation or thought came across me, which was opposed to my general results, to make a memorandum of it without fail and at once; for I had found by experience that such facts and thoughts were far more apt to escape from the memory than favourable ones."

$$\text{s.t. } \mu_{\tau'}(s_{\tau'}) = \mu_{\tau'}^B(s_{\tau'}) \quad \forall \tau' > t, \forall s_{\tau'}.$$

The analysis of agents is according to their levels of introspection and patience. I start by considering agents that are highly sophisticated: forward looking agents. Forward looking agents realize the effects of their manipulations of beliefs at period t on their belief formation at period $t + 1$, and their utility consequences. I then contrast the results of the forward looking agents with the analysis of the general case, agents who have imperfect perceptions of their future behavior.

D. Belief Utility

Psychological studies seem to suggest that people have a taste for confidence (see Griffin and Tversky [1992] for many references of excessive confidence in many realms of scientific and professional life). I consider a special case of belief utility which will be termed *directional confidence*. This belief utility captures the idea that the agent has a taste for having beliefs that support the same action as did previous beliefs. Moreover, the agent likes to have strong convictions. In other words, the agent welcomes more confidence that confirms her priors.

Formally, I use a linear specification.⁷ The belief term of the utility at stage $t + 1$ is given by:

$$v(\mu_t, \mu_{t+1}) = \begin{cases} \mu_{t+1} - \mu_t & \mu_t > \frac{1}{2} \\ |\mu_{t+1} - \mu_t| & \mu_t = \frac{1}{2} \\ \mu_t - \mu_{t+1} & \mu_t < \frac{1}{2} \end{cases}.$$

Thus, the agent gains belief utility if she becomes stronger in her convictions that her previous period's action was the optimal one.⁸

⁷A more general form of directional confidence would be:

$$v(\mu_{t+1}, \mu_t) = \begin{cases} f_L(\mu_{t+1} - \mu_t) & \mu_t > \frac{1}{2} \\ f_I(\mu_{t+1} - \mu_t) & \mu_t = \frac{1}{2} \\ f_R(\mu_{t+1} - \mu_t) & \mu_t < \frac{1}{2} \end{cases}$$

where $f_L(0) = f_I(0) = f_R(0) = 0$ and $\frac{\partial f_L}{\partial x}(x) > 0$, $\frac{\partial f_R}{\partial x}(x) < 0$, $\frac{\partial f_I}{\partial x}(x) > 0$ for $x > 0$ and $\frac{\partial f_L}{\partial x}(x) < 0$ for $x < 0$. Analysis of this form of the belief utility yields similar qualitative results, but is far more cumbersome.

⁸Another prominent consistency story in the psychology literature is that of disappointment and perceived utility maintenance (see, for example, Tesser [1988]). An agent at time t suffers disappointment if current perceived utility from her past actions (as calculated using her adopted beliefs) is lower than her period $t - 1$ perceived utility. Define:

$$v(\mu_1, \mu_0) = \text{dissapointment}_1(a_0, \mu_1) = w((\mu_1 - \mu_0)I_L(a_0) + (\mu_0 - \mu_1)I_R(a_0))$$

where $I_j(k) = \begin{cases} 1 & j = k \\ 0 & j \neq k \end{cases}$ and a_0 is assumed to be maximize the instantaneous utility with first stage beliefs μ_0 .

Taking $w(x) = x$, we get:

Remark 3: I call an alternative form of belief utilities *regret*. Assume that for any belief μ_{t+1} the agent chooses, she calculates her expected past earnings from her choice of stage t action a_t and the maximum expected value she could have gotten had she used μ_{t+1} in order to choose past actions. If the latter is greater than the former, the agent experiences regret.

Formally, we assume that the belief term takes the form:

$$v(\mu_{t+1}, \mu_t, a_t) = R(\mu_{t+1}I_L(a_t) + (1 - \mu_{t+1})I_R(a_t) - \max\{\mu_{t+1}, 1 - \mu_{t+1}\}) =$$

$$= \begin{cases} 0 & a_t = L, \mu_{t+1} \geq \frac{1}{2} \\ 1 - 2\mu_{t+1} & a_t = R, \mu_{t+1} \geq \frac{1}{2} \\ 2\mu_{t+1} - 1 & a_t = L, \mu_{t+1} \leq \frac{1}{2} \\ 0 & a_t = R, \mu_{t+1} \leq \frac{1}{2} \end{cases}$$

when the regret function R is taken to be linear. Here previous beliefs play a role in the belief utility only to the extent that they determine the previous period's action. It is worth noting that regret functions do appear in the economics literature. For references and some theoretical analysis of regret functions, the reader is directed to Loomes and Sugden [1982]. In this paper we present results for directional confidence belief utilities. Regret and joy utilities lead to similar qualitative results and are not analyzed here.

IV. Main Results

The main results are of two types. The first refers to the kind of beliefs the agent adopts. I give conditions for the agent being overconfident and underconfident. These results are summarized in Subsection A. Since actions are strongly tied to the beliefs the agent holds, a taste for consistency may have observable implications on the actions choices. Characterization of conditions that lead to persistence of choices and apparent habit formation are presented in Subsection B. The second type of results pertains to the levels of expected utility the agent will experience, as is perceived by the agent, and by an omniscient entity (or a plain Bayesian), at the outset of the process, as a function of the amount and quality of the information provided. These results are summarized in Subsection C. The proofs of the Propositions stated in this section are relegated to Appendix A.

$$v(\mu_1, \mu_0) = \begin{cases} \mu_1 - \mu_0 & \mu_0 > \frac{1}{2} \\ \mu_0 - \mu_1 & \mu_0 < \frac{1}{2} \end{cases}.$$

For $\mu_0 = \frac{1}{2}$ we can arbitrarily choose $v(\mu_1, \mu_0) = |\mu_1 - \mu_0|$ and v is then identical to the linear directional confidence belief.

A. Choice of Beliefs

When agents have a taste for consistency as suggested in our framework, next period's signal will have an effect on the current period's beliefs for a non-myopic agent. In particular, beliefs at period 1 may depend on the accuracy of the future signal q_2 . While this does not occur in standard learning models (e.g., if we take $\gamma = \tilde{\gamma} = 0$ in the present model), it is quite intuitive in the current setup. Indeed, q_2 , together with $\tilde{\gamma}$ determine the agent's forecast concerning the extent to which period 2's beliefs will be modified. For example, if $\tilde{\gamma}$ is very small and q_2 is very high, the agent predicts that in period 2 she will not care about consistency much, hence be using Bayes' rule. Since q_2 is high, the agent thinks she will be guessing the right state of the world with a high probability. Hence, she can afford to distort her current beliefs and gain the additional confidence in period 1. This line of arguments would not hold if q_2 were very low, in which case the agent would predict very little future changes to her beliefs, and may choose period 1's beliefs more conservatively.

If q_2 is very high, then for sufficiently low $\tilde{\gamma}$ the agent predicts that she will be following Bayesian updating in the future and hence forecasts that her second period's beliefs will be very accurate. Thus, the agent is more likely to distort her current beliefs. If $\tilde{\gamma}$ is very high the current modification is perceived likely to have a significant impact on the future belief. In that case, since the agent wants to give leverage for updating when q_2 is very high, she will be more likely to choose less extreme beliefs. Hence, beliefs at stage 1 are likely to be affected by the accuracy of the second stage payoff.

The agent is said to be overconfident (underconfident) relative to a Bayesian observer if her beliefs are too extreme (too conservative) relative to the Bayesian posterior. Formally,

Definition (underconfidence and overconfidence): An agent holding a belief μ is overconfident relative to a belief μ^B if:

1. $\mu > \mu^B$ and $\mu^B > \frac{1}{2}$; or
2. $\mu < \mu^B$ and $\mu^B < \frac{1}{2}$.

Analogously, an agent holding a belief μ is underconfident relative to a belief μ^B if:

1. $\mu < \mu^B$ and $\mu^B > \frac{1}{2}$; or
2. $\mu > \mu^B$ and $\mu^B < \frac{1}{2}$.

I will henceforth refer to underconfidence and overconfidence only relative to the beliefs a Bayesian observer would construct.

When the accuracy of the first signal is sufficiently high and the agent cares about her beliefs, she will distort first period beliefs. Moreover, the distortion would be in the direction that increases belief utility. Since the prior is favorable to $\theta = L$ ($p > \frac{1}{2}$), the agent will choose a belief above the posterior when guessing L in the first period (increase her confidence) and choose a belief that is above the posterior when guessing R (since in this case she loses by her inconsistency with the

prior, and would like to minimize the loss by choosing a belief that is closest to the prior p). In such a situation, the agent would appear either overconfident (when $s_1 = L$) or underconfident (when $s_1 = R$).

If beliefs in the first period are chosen extreme enough, then beliefs in the second period may still be too extreme relative to the true Bayesian posterior based on both signals. As it turns out, this situation arises when agents care enough about consistency, in the sense of a sufficiently high parameter γ^* , and have a sufficiently accurate first period signal. Formally,

Proposition 1 (exaggerated confidence levels): *There exists a $\gamma^* \in (0, 1]$ such that for all $\gamma \geq \gamma^*$ there exists $q_1^*(\gamma) \in (0, 1]$, $q_1^*(\gamma) < 1$ when $p < 1$, such that for all $q_1 > q_1^*(\gamma)$, for all $q_2 \in [\frac{1}{2}, 1)$, the agent is overconfident when $s_1 = L$ and underconfident when $s_1 = R$ in both periods.*

Intuitively, a small change of first period beliefs within either the range $(0, \frac{1}{2})$ or $(\frac{1}{2}, 1)$ doesn't affect the actions chosen. Hence, first period beliefs balance the tradeoff between getting the belief utility in the first period and probabilistically losing through the instrumental utility in the second period. A marginal increase in the level of beliefs in the first period changes the level of utility by that same increase multiplied by γ , and yet possibly reduces the second period utility through the instrumental part by the probability of guessing wrong because of insufficient updating. When γ is high, and the accuracy of the first signal is sufficiently high, the first effect dominates the latter.

B. Choice of Actions

Since actions are strongly tied with beliefs, persistent beliefs will be connected with persistent actions.

When the agent perceives her future self not to care much about beliefs, she will predict that her behavior in the future will be approximately Bayesian, and hence will allow herself to choose more extreme beliefs. However, when stage 2 comes to be, the agent does care about consistency, and if first period beliefs are chosen extreme enough, will not be willing to change her mind. In such a situation, the agent will appear to have sticky choices and first period information will determine the actions throughout the process with no regard to the additional information that arrives at period 2. This intuition is presented in the following proposition.

Proposition 2 (history matters): *There exists a $\gamma^* \in (0, 1]$ such that for all $\gamma \geq \gamma^*$ there exists $q_1^*(\gamma)$ such that for all $q_1 \geq q_1^*(\gamma)$, for all $q_2 \in [\frac{1}{2}, 1)$, $q_2 > \tilde{\gamma}$, $s_1 = L \Rightarrow a_1 = a_2 = L$ no matter what s_2 is, and yet $s_1 = R \Rightarrow a_1 = a_2 = R$ no matter what s_2 is.*

It is important to note that the conditions of Proposition 2 allow, in particular, for $q_2 > q_1$, in

which case with positive probability the agent is in fact choosing an action in the second period that a standard agent (with $\gamma = \tilde{\gamma} = 0$) would not choose. In particular, with positive probability the agent is not achieving the maximal instrumental utility her information would potentially allow her to achieve.

Another way to describe the significance of the history of actions and beliefs is as follows. So far, I've assumed symmetry between the two states L and R : guessing correctly that $\theta = L$ leads to the same instantaneous utility as guessing correctly that $\theta = R$. If we allow for a difference in tastes between guessing $\theta = L$ correctly and guessing $\theta = R$ correctly, the utility can be written as:

$$u(a_t) = \begin{cases} \lambda & a_t = \theta = L \\ 1 - \lambda & a_t = \theta = R \\ 0 & otherwise \end{cases} .$$

If, e.g., $\lambda > \frac{1}{2}$ this means that the agent would prefer to guess $\theta = L$ correctly over guessing correctly that $\theta = R$. It is as if the agent prefers the state L over R .

Let $I_x(y)$ denote the indicator function. That is,

$$I_x(y) = \begin{cases} 1 & x = y \\ 0 & otherwise \end{cases} .$$

Then, one can write:

$$u(a_t) = \lambda I_L(a_t) I_L(\theta) + (1 - \lambda) I_R(a_t) I_R(\theta).$$

Suppose an econometrician observes the following:

- a. The agents' prior concerning the state of the world μ_0 ;
- b. The signals $\{s_t\}$ the agents observe; and
- c. The actions $\{a_t\}$ the agents have picked.

Assume that at the last stage T the (Bayesian) posterior is calculated to be μ_t^B and the action the agent takes is $a_t = L$. Then if the agent maximizes a utility function of the above form, she should have calculated that

$$\begin{aligned} E_{\mu_T^B}(u(L)) &\geq E_{\mu_T^B}(u(R)) \Leftrightarrow \\ \lambda \mu_T^B &\geq (1 - \lambda)(1 - \mu_T^B) \Leftrightarrow \lambda \geq 1 - \mu_T^B. \end{aligned}$$

Hence, the econometrician would be able to *estimate* a threshold λ for each agent.

Proposition 2 suggests that, in a two period scenario, for a certain range of parameters, the

estimated λ of the second period would be correlated with the first period's signal. The agent would seem to "like" L more than R if L is the guess she made at period 1, and vice versa. This is reminiscent of the phenomena of habit formation in which agents' actions become repetitive, or habitual, beyond what Bayesian information processing would justify.

Formally, we have the following corollary:

Corollary (apparent habituation): *There exists a $\gamma^* \in (0, 1]$ such that for all $\gamma \geq \gamma^*$ there exists $q_1^*(\gamma)$ s.t. for all $q_1 \geq q_1^*(\gamma)$, for all $q_2 \in [\frac{1}{2}, 1)$, $q_2 > \tilde{\gamma}$, if a Bayesian statistician observes p, s_1, s_2, a_1 , and a_2 and hypothesizes that the agent's objective function, at least at the last period, is of the form:*

$$u(a_t) = \lambda I_L(a_t) I_L(\theta) + (1 - \lambda) I_R(a_t) I_R(\theta)$$

then she will deduce that when $s_1 = L, s_2 = R$, $\lambda \geq \lambda_L^$ and when $s_1 = R, s_2 = L$, $\lambda \geq \lambda_R^*$ where $\lambda_L^* > \lambda_R^*$.*

C. Information Quality and Expected Utility

I now turn to analyze the effects of information quality on the agent's expected utility. When the agent does not care about consistency, as is commonly assumed, more accurate information is always weakly preferred to less accurate information. The reason is that information can always be ignored. Hence, an agent with an accurate signal can imitate an agent with a less accurate signal, but not vice versa.

In the present setting, there are two natural approaches to the comparative statics analysis related to information accuracy. The first is paternalistic - knowing how the agent actually behaves (in particular, knowing the agent's propensity to behave inconsistently), one can determine which structures of information would lead the agent to higher expected utility. I use the term *ex-ante expected utility* to describe the total expected utility of the agent before any signal is realized, taking into account her actual behavior in all periods. That is, taking into account the agent's possible lack of consistency.

Another approach deals with the agent's aspect. In a market situation, the agent might need to choose which signals to purchase or pay attention to by herself. Hence, one can study which information structure the agent would prefer at the outset of the game. I use the term *perceived expected utility* for the utility the agent expects to achieve before any signal is realized (given her level of introspection, as captured by γ and $\tilde{\gamma}$).

As it turns out, both the ex-ante expected utility and the perceived expected utility (at stage 0) are not monotonic in the accuracies of the signals.

I start by looking at the ex-ante expected utility. Intuitively, assume that $\tilde{\gamma}$ is very small and that γ is significant, so that the agent cares about consistency in period 1 but believes she will care very little in period 2, so that she will use the Bayesian update in the future. If current accuracy is very high, the agent might be willing to take extreme beliefs, counting on the fact that approximate Bayesian updating will occur in the second period. However, since the agent actually does care about consistency of beliefs in the second period, she might not update her beliefs when first period beliefs are extreme in order to avoid a loss through the belief term. Thus, there are two contradicting forces. On the one hand, higher first period accuracy would imply higher expected instrumental and belief utilities in the first period. On the other hand, the resulting extreme beliefs in the first period might cause a reduction in instrumental utility in the second period. For certain parameters, the latter effect is greater than the former, and we get non-monotonicity with respect to the first period accuracy q_1 .

In a similar spirit, if the future accuracy is very high, the agent expects to bear a relatively low future instrumental cost when choosing an extreme belief - future signals will be so accurate that they will lead to the right choice no matter what current belief is chosen. The agent might be more overconfident when future accuracy of signals is higher. However, when stage 2 arrives, the agent does care about consistency between her beliefs and might end up exhibiting a relatively low willingness to change her first period beliefs, thereby losing some instrumental utility. When the signal of the second period is less accurate, the beliefs of stage 1 are chosen more modestly and hence are more likely to be modified in the second period. The implication is that information may be better utilized, from a Bayesian outlook, when the signal in the second period is lower. This intuition is captured in the following proposition.

Proposition 3 (ex-ante utility non-monotonicity): 1. (w.r.t. q_1) There exist $\delta^*, \gamma^*, \tilde{\gamma}^*$, such that for all $p < \frac{1}{\sqrt{2}}, \delta > \delta^*, \gamma > \gamma^*$, and $\gamma > \tilde{\gamma} < \tilde{\gamma}^*$ there exist $q_1^L, q_1^H, q_2 \in (\frac{1}{2}, 1), q_1^L < q_1^H$ s.t. the ex-ante expected utility with signal accuracies q_1^L and q_2 is higher than that with accuracies q_1^H and q_2 .

2. (w.r.t. q_2) There exist $q_1^*, \delta^*, \gamma^*, \tilde{\gamma}^*$, such that for all $q_1 > q_1^*, \delta > \delta^*, \gamma \geq \gamma^*, \gamma \geq \tilde{\gamma} \leq \tilde{\gamma}^*$, ex-ante expected utility is non-monotonic in q_2 .

I now turn to the analysis of perceived expected utility. The agent can always ignore or distort information in the first period, hence more accurate information in the first period cannot decrease her perceived well-being. Nonetheless, in our setup, the agent predicts her preferences will change in the future, namely that she will care less about consistency between beliefs than she does at the outset of the process. Since the agent cannot commit herself to take a certain future action, there is an intrinsic agency problem, and the agent may prefer less accurate signals for the future.

To make this intuition sharper, assume, as before, that $\tilde{\gamma}$ is very small and that γ is significant, so that the agent cares about consistency in period 1 but believes she will care very little in period 2, so that she will use the Bayesian update in the second period. In period 1, the agent cares about the changes in beliefs both in period 1 and in period 2. When the accuracy of the signal in the second period is very high the agent thinks, at period 1, that she is more likely to change her beliefs in the second period, eventhough from period 1's perspective, she would prefer to be more conservative. Therefore, there are two opposing forces at play. On the one hand, more accurate signals lead to higher instrumental utility. On the other hand, the agent perceives her future self to ignore her current taste for consistency, and might perceive a greater loss on the overall utility in the second period when future signals are very accurate. This effect arises from the agent's perceived inability to enforce her current preferences on her future self. For certain parameters, the second effect dominates the first and perceived utility is non-monotonic in the accuracy of the second signal. More formally,

Proposition 4 (perceived utility (non-)monotonicity): *1. For all parameters, perceived utility is non-decreasing in q_1 .*

2. There exist δ^ and γ^* such that for all $\delta > \delta^*$ and $\gamma > \gamma^*$ there exist $\tilde{\gamma}, q_1, q_2^L, q_2^H, q_2^L < \tilde{\gamma} < q_2^H$ such that perceived expected utility with $(\delta, \gamma, \tilde{\gamma}, q_1, q_2^L)$ is higher than that with parameters $(\delta, \gamma, \tilde{\gamma}, q_1, q_2^H)$.*

In the last two propositions, we were concerned with the *quality* of information the agent would prefer. A related question concerns the *quantity* of information the agent would choose. I compare a one stage process with a two-stage one where the first period signal is of identical accuracy. I.e., in the two-stage scenario, there is additional information in period 2. In order to make meaningful comparisons between the two, I normalize the two-stage game's payoff by $\frac{1}{1+\delta}$.

Intuitively, it is useful to consider the case of $\delta = 1$. In this case, the belief utilities of stages 1 and 2 receive a weight of $\frac{1}{2}$. When the accuracy of the second signal, q_2 , is low, the agent's action will remain the same across periods. Hence, the agent moves from a prior p to a posterior of at most 1, leading to an addition of $\frac{1-p}{2}$ to the belief term in the utility, or a posterior of at least 0, leading to a maximal addition of $\frac{1}{2}[|0 - \frac{1}{2}| + (\frac{1}{2} - p)] = \frac{1-p}{2}$ to the belief term. If q_1 is sufficiently low, the agent will always become overconfident in the one-stage game and adopt belief $\mu_1 = 1$, leading to an additional $(1 - p)$ utils to the belief term. Hence, in the two-stage scenario the agent gains from the additional confidence only half as much as in the one stage setup. It turns out that for certain parameters, the instrumental part of the utility does not counter this effect. In these situations, the agent strictly prefers the one-stage scenario, despite the fact that she receives less information then. The interpretation of this is that an agent that has a taste for increasing confidence per-se might

forgo information at the expense of instrumental utility. The following proposition formalizes this intuition.

Proposition 5 (preferring less information): *There exist $\delta^*, \gamma^* \in (0, 1)$ such that for all $\delta > \delta^*$, $\gamma > \gamma^*$, there exists a q_2^* such that, for all $q_2 < q_2^*$, all $q_1 < \min\{\frac{3p}{1+2p}, \frac{1+p}{2}\}$ the agent strictly prefers the one-stage game with stage 1 signal of accuracy q_1 over the two stage-game with signals s_1 and s_2 of accuracies q_1 and q_2 , respectively, when payoffs of the two-stage game are normalized by $\frac{1}{1+\delta}$.*

Remark 4: In this model uninformative signals, or signals that don't affect the agent's choice of actions may increase the agent's perceived expected utility. Indeed, even in a one-stage scenario, as long as the accuracy of the signal is less than the prior p , the agent would stick to the action determined by the prior (in our framework, the action L), but become more confident and gain "belief utility." In fact, as long as the price of the signal is less than $\gamma(1 - p)$, the resulting gain from the increased confidence, the agent would be willing to purchase such a signal.

Remark 5: Related to the previous remark, given that sometimes uninformative signals serve to increase the expected utility of an agent who cares about her beliefs (characterized by $\gamma > 0$), one might wonder why any event cannot serve as a signal for such an agent. A few possible answers can be given:

1. As Kunda [1990] notes, people are indeed more likely to believe things they want to believe, but their capacity to do so is constrained by objective evidence and by their ability "...to construct a justification of their desired conclusion that would persuade a dispassionate observer. They draw the desired conclusion only if they can muster up the evidence to support it." Hence, people cannot use absolutely any event to justify their hypotheses.
2. Related to 1, signals seem to be multi-dimensional. We could think of a (binary) dimension that indicates the relevance of each signal to the issue at hand. Signals may be used for analysis only if they are relevant. The analysis presented in this paper is then for signals that are relevant for the decision problem at hand.
3. Alternatively, we could think that signals can be modified only to a certain extent which depends on their accuracy. That is, upon constructing a belief using Bayes' rule, the agent can choose her beliefs within a range around the "correct" belief, the size of this range being proportional to the accuracy of the signal. In particular, beliefs based on non-informative signals cannot be modified. Formalizing this idea does not change the analysis qualitatively, but makes it quantitatively much more cumbersome.

Remark 6: As it turns out, the results for forward looking agents, agents for whom $\tilde{\gamma} = \gamma$ and who are hence time consistent, are somewhat special. A forward looking agent understands fully her future actions and can thus use backwards induction arguments correctly. In this case “free disposal of information” always holds, in the sense that information can always be ignored or distorted. In particular, more accurate signals are always weakly preferred to less accurate ones, and additional signals never decrease expected ex-ante utility. Some of these results are presented in Appendix B.

D. So what are γ and $\tilde{\gamma}$ really?

Empirical studies calibrating the numerical levels of γ and $\tilde{\gamma}$ are still lacking. However, a recently growing literature in social psychology deals with what is termed *affective forecasting*, the prediction of one’s feelings and emotions. Studies show that people are not very good forecastors of their future reactions to events. For example, Gilbert et al. [1998] and Loewenstein and Schkade [1999] ran numerous studies asking people who did and did not go through an unpleasant event (such as divorce, denial of tenure, sudden disability, etc.) for their reactions. Their results clearly illustrate a discrepancy between people’s predictions concerning their reactions and people’s actual reactions to these events. Loewenstein et al. [2000] examined people’s attitudes toward risk (real time as opposed to distant event) and discovered a similar bias.

These studies imply that while γ , the parameter that symbolizes the weight that is given to being consistent, might be substantial, people cannot introspect perfectly and predict γ accurately. That is, qualitatively, it seems quite plausible that $\tilde{\gamma} < \gamma$.

Morgan et al. [1998] study the relationship between patient coping style (taste for information), precolonoscopy information, and anxiety and pain associated with colonoscopy. Based on the Krantz Health Opinion Survey (Krantz et al. [1980]), the authors can classify their subjects into two groups: information seekers and information avoiders. Patients given information congruent with their coping style experienced significantly less self-reported anxiety immediately after the information intervention and spent less time in recovery. These kinds of results hint that people differ in the parameters that characterize their behavior. Moreover, such studies could potentially serve to help calibrate γ .

E. Is it just the self?

One might suggest that the phenomena described and modeled in this paper are driven from people’s desire for self-enhancement. That is, people want to believe they are able, in particular, that they take the right actions and hold the right theories, and information distortion occurs in order to maintain, or even enhance, the notion of competence. These ideas were formally put

forward by Akerlof and Kranton [2000] and Koszegi [1999].

In the social psychology literature, there has been an ongoing debate between two main schools. The first poses that people act according to self-enhancing motives (see Baumeister [1998], Tesser [1988], and references therein), that indeed people tend to interpret events in a way that enhances their self-concept. The second school of thought supports the notion of self-verification, a term coined by Swann [1985]. This school suggests that people have a taste for consistency in their self-perceptions, even if those are not good. In a long series of experiments, Swann and his colleagues tried to exhibit people's quest for feedback that confirms their views about themselves. For example, Swann et al. [1994] have studied nearly 200 couples and discovered that married folks with a negative self-image are more intimate with spouses who evaluate them unfavorably than with partners who lavish them with seemingly undeserved praise. Even those with a positive self-view may psychologically withdraw from a marriage if their mate seems unjustifiably effusive. In another experiment Pelham and Swann [1989] have shown that people tend to be particularly receptive to negative personal feedback when they suffer from low self-esteem.

I choose not to enter this debate and therefore I do not attempt to connect the results presented here with ideas about signaling of competence (to oneself or to others).

V. Applications

Overconfidence has been considered in several economic domains. In finance, Shleifer [2000] proposes illustrations of overconfidence of financial investors. Akerlof and Dickens [1982] provide evidence for overconfidence of workers employed in hazardous occupations. In the legal world, experiments have shown that judges become more and more confident in their position, even when they are given evidence that is irrelevant to the case at hand (see Nisbett and Ross [1991]).

I choose to sketch a few applications of the results on excess stickiness of actions and choice of information structure. I consider applications to political campaigns as well as marketing and business strategies.

A. Political Campaigns

I draw the following analogy. Each voter holds a prior belief as to which, and by how much, a candidate or policy is better. Campaigns, debates, media reports, etc. serve as signals the agent may or may not receive.

It is assumed that listening to debates or to the campaigns of one's opposition may bring new arguments to mind and is thus more informative than listening to campaigns that replicate one's opinions. Hence, agents have a choice of both the quantity and the quality of signals they get exposed to.

Biases in the composition of voluntary audiences to mass communications have been reported often in survey studies (commonly referred to as *selective exposure*). Lazarfeld et al. [1948] were some of the first to point out that “most people expose themselves, most of the time, to the kind of propaganda with which they agree to begin with.” One of the classical findings they reported is that of respondents with constant voting intentions from May to October of 1948. About two thirds were exposed predominantly to propaganda favoring their side, and less than one fourth mainly to propaganda favoring the other side. That is, people do not choose the most accurate information available to them.

Sears and Freedman [1967] mention the example of Senator William Knowland’s telethon in the 1958 California gubernatorial election. Interviews with voters immediately after the election revealed that:

1. Twice as many Republicans as Democrats (proportionately) had seen this Republican candidate’s program.
2. The average Republican viewer watched the program for about an hour longer than did the average Democratic voter.

These observations are consistent with the results on agents’ preference for information that is not maximally accurate.

More recently, Ansolabehere and Iyengar [1994] use an experimental setting to show that candidates gain most from advertising on issues over which they can claim “ownership,” issues over which the candidates have a preconceived advantage in the electorate (for example, the public generally considers Democrats to be more able than Republicans in solving problems of unemployment and civil rights). In the experimental setting, subjects did not have a choice of whether or not to listen to an (orchestrated) advertisement. Rather, the authors tested the subjects’ interpretation of different kinds of advertisements.

B. Marketing Strategies

Once agents do not perfectly realize their future taste for consistency, there is room for exploitation by an outside agent. The principle of “foot in the door” refers to a situation in which asking for a small favor or commitment (one that, in the context at hand, can hardly be refused), and only then asking for the larger commitment raises the prospects of the latter being undertaken. Freedman and Fraser [1966] demonstrated this phenomenon in a classical experiment. Homemakers in a middle-class housing tract near the Stanford University campus were first approached by a person who asked them to do something relatively innocuous (sign a petition or put a small sign in the window of their car or home promoting a noncontroversial cause, like safe driving). The vast

majority of those approached agreed.⁹ Two weeks later a second person visited the same sample of homemakers, and also called upon a control group sample who had not been previously contacted, to accede to a far more substantial, even rather unreasonable, request. He asked them to allow a large, crudely lettered, and decidedly ugly “Drive Carefully” sign to be installed directly in front of their house. As he made this request, he showed them a photo in which the ugly sign could be seen obscuring the front door of another house in the tract.

Fully 76 percent of the subjects who had initially agreed to place a small auto safety sign in their window now agreed to place the big and ugly “Drive Carefully” sign in the front of their house. By contrast, “only” 17 percent agreed to erect the sign when there had been no prior foot in the door visit.

A similar study by Pliner, Hart, Kohl, and Saari [1974] found that a sample of Toronto suburbanites became twice as likely to donate money to the Cancer Society after agreeing, a day earlier, to wear a lapel pin publicizing the forthcoming fund drive (a request that none of the subjects actually refused).

A related method for luring consumers is that of deceptive advertisement. If consumers are convinced to purchase a product in period 1, a taste for consistency will push them to appreciate the product even if the advertisement that led them to make the purchase was false. Hence, firms may have an incentive to advertise deceptively.¹⁰

This evidence is related to my results on excess persistence of actions. The “foot in the door” experiments do not fit the framework perfectly since the benefits from complying with the requests in different stages of the process seem to change. In order to use the suggested framework one would need to generalize my results to a situation in which u depends on t . For example, one could assume that $u(a_0 = \theta) = u(a_1 = \theta) = 1$, $u(a_2 = \theta) = \eta (< 1)$ while $u(a_t \neq \theta) = 0$ for all t . Similar results to Proposition 2 hold.

C. Business Decisions

In the organizational behavior literature cognitive dissonance is sometimes termed *escalation of commitment*. It refers to situations in which once people choose one course of action, their beliefs change in a way that makes them committed to their chosen course. This common wisdom corresponds to the persistence results presented in Proposition 2. One of the first studies in the organizational behavior literature of this sort was done by Staw [1976]. Staw let subjects (business school students) play the role of corporate financial officer in allocating R&D funds to the operating

⁹Unfortunately, the authors do not specify the percentages of compliance.

¹⁰Nagler [1993] provides some evidence for deceptive advertising. Nagler also presents a specific model of cognitive dissonance in which a fraction of consumers incurs a switching cost with respect to actions. The model shows how firms may benefit by advertising deceptively.

divisions of a hypothetical company. Half the subjects allocated were given feedback on their decision and were then asked to make a second allocation. The other half of the subjects did not make the decision themselves, but were told that it was made by another financial officer of the firm. Feedback was manipulated so that half the subjects received positive results on their initial decisions, while half received negative results.

The data showed that subjects allocated significantly more money to failing than to successful divisions. It was also found that more money was invested in the chosen division when the participants, rather than another financial officer, were responsible for the earlier funding decision. Unfortunately, the author did not question subjects for their beliefs on the future success of the different divisions.

One would wonder whether business managers also collect information that is not maximally accurate, as predicted by Proposition 4. Unfortunately, there is not much evidence on this question. One related report is by Ehrlich et al. [1957] who found that people, whether or not they had just bought new cars, read a higher percentage of the available ads about their own makes than about any other make.

VI. Strategic interactions

In this section I suggest some illustrative results on learning processes in games between players who have a taste for consistency.

Consider the following game (matching pennies):

	Left	Right
Top	1, -1	0, 0
Bottom	0, 0	2, -2

and consider players who are playing fictitious play with a taste for beliefs as well. That is, Player 1, the row player, starts with an initial belief, which is the probability associated to Player 2 playing L, of:

$$\mu_0 = \frac{N_0^L}{N_0^L + N_0^R}$$

and the fictitious play belief after t stages is:

$$\mu_t^{FP} = \frac{N_0^L + N_t^L}{N_0^L + N_0^R + t}$$

where N_t^L is the number of times the action L has been played in the first t periods of the game.

At each stage, when choosing a belief μ_t , the agent has a trade-off between the instrumental effect of $\mu_t^{FP} - \mu_t$ and the belief effect which depends on $\mu_t - \mu_{t-1}$.

The action chosen at time t , a_t , is given by:

$$a_t = \begin{cases} T & \mu_t > \frac{1}{2} \\ B & \mu_t < \frac{1}{2} \end{cases}$$

Hence, for a myopic agent, for sufficiently high μ_0 , $a_t = R \forall t$.

Proposition 6: *Let $A = (a_{ij})$ be a zero-sum game, a_{ij} denoting the ij payoff to Player 1. Let I^* denote the set of actions which are not weakly dominated. Assume Player 1 is myopic and Player 2 is sophisticated, and both use fictitious play dynamics. Then there exists a set of initial beliefs of Player 1 with which Player 2 can guarantee a steady state payoff of $\min_{i \in I^*, j \in J} a_{ij}$.*

Remark 7: For a sufficiently forward looking Player 2, there is no restriction on the initial beliefs of Player 1 since Player 2 can play a certain action for an initial block and drive Player 1 to any belief.

When there are more than two actions for either player, beliefs are more than one-dimensional and we need to modify the functional form of the belief utility. For ease of exposition, I will consider only forms that are extensions of the taste for directional confidence. Once each action is associated with a number (I will assume different actions are associated with different numbers), one can calculate the variance of a distribution over actions. If μ is a belief over actions, $var\mu$ is the variance of the corresponding random variable. Since confidence is associated with the variance of the distribution, a natural generalization of the belief utility for one-dimensional beliefs is of the form:

$$v(\mu_t, \mu_{t-1}) = \begin{cases} \|\mu_t - \mu_{t-1}\| & var\mu_t < var\mu_{t-1} \\ \|\mu_t - \mu_{t-1}\| & var\mu_t = var\mu_{t-1} \\ -\|\mu_t - \mu_{t-1}\| & var\mu_t > var\mu_{t-1} \end{cases}$$

where $\|\cdot\|$ denotes the Euclidean norm.

Proposition 7 (extension of Proposition 6 to non-zero-sum games): *Let $(A, B) = ((a_{ij}), (b_{ij}))$ denote a normal form two-player game. Assume Player 1 is myopic and Player 2 is sophisticated, and both use fictitious play dynamics. Let I^* denote the set of actions of Player 1 which are not weakly dominated. That is, each action in I^* is a best response against some*

action of Player 2. Then there exists a set of initial beliefs of Player 1 with which Player 2 can guarantee a steady state payoff of

$$\max_{i \in I^*, j \in J} b_{ij}.$$

Proof: Let $b_{ij} \in \arg \min_{i \in I^*, j \in J} b_{ij}$. From the definition of I^* there exists an action $j(i)$ such that $i \in BR(j(i))$. Since there is a finite number of actions of Player 2, there is a set of Player 1 beliefs Ω with positive (Borel) measure against which i is a best response. Let $\mu = (\mu_1, \dots, \mu_J) \in \Omega$ (by a slight abuse of notation I assume that the number of Player 2's actions, $|J|$, is equal to J).

Pick $N_0^k = 0 \ \forall k \neq j(i)$, $N_0^{j(i)} = N^*$ s.t. $i \in BR(\mu')$ for all μ' of the form $\mu'_k = \frac{1}{N+1}$ for some $k \neq j(i)$ and $\mu'_{j(i)} = \frac{N}{N+1}$ for all $N > N^*$ (there exists such an N^* since J is finite).

Note that for such μ' , $var \mu' > var \mu_0$ (no matter how the actions were numbered).

Hence, whatever action is taken at stage 1 of the game it is optimal for Player 1 to choose $\mu_1 = \mu_0$. Inductively, $\mu_t = \mu_0$ for all t and hence Player 1 will play i for all t when the initial belief is μ_0 . ■

Remark 8: In the standard fictitious play, players hold indices for the number of times each action has been played. Here, we have been using a less restrictive assumption on players' sophistication. Each player remembers at each stage t the pair $(\mu_{t-1}, N_0^1 + N_0^2 + \dots + N_0^J + t)$ where J is the opponent's number of actions.

VII. Concluding Comments

The psychology literature indicates that people exhibit cognitive biases implying a taste for consistency. In broad terms, this paper's contribution is in providing a framework for analyzing decision making when agents choose actions and beliefs. In particular, the paper proposes a model in which distortions of beliefs from the Bayesian posteriors arise endogenously.

The framework helps explain under- and overconfidence, excess stickiness in action choices, and agents' preference for less accurate signals or less signals altogether. The results are consistent with evidence concerning observed overconfidence of financial investors, selective exposure to political information, marketing strategies exploiting people's tendency to persist with a taken action, escalation of commitment of business managers, and more.

The underlying setup suggested in this paper offers grounds for more work, theoretical and applied. In what follows I suggest a few directions for future inquiry.

On a theoretical level, the model presented here is a benchmark model in a few respects. For example, I assumed the agent has only one stage memory of beliefs. There is some evidence indicating that people tend to remember their theories, rather than the hard evidence that led to their construction (see Shacter [1996]). However, a possible extension of the current model would look at longer decision processes and allow for the memorization of more than one belief. In order to have a realistic model, one would need to take into account effects of selective memory, particularly, the tendency to remember things that enhance one's sense of well-being. In other words, what people remember may be partly endogenous to the structure of their preferences as well.

Another strong assumption I posed refers to agents maximizing their objective function with respect to the Bayesian posterior, that may differ from the belief chosen that period. One could potentially relax this assumption in several ways - limiting the extent to which the adopted belief can differ from the Bayesian posterior, letting the agent maximize her objective function with respect to some convex combination of the Bayesian update and her adopted belief (see footnote 4), and the like.

The results on agents' preference for less accurate information over more accurate information serve as grounds for potential new experiments on information acquisition, as well as a natural point for testing the theory suggested in this paper. In addition, such results would suggest that the pricing of information in the market would not necessarily be monotonic in the amount of information. The model would suggest a preference relation over information structures that would be qualitatively different from, e.g. the commonly used Blackwell relation (see Blackwell [1950]). A reasonable cost function would then be a representation of this relation.

In Section VI, I sketched preliminary results related to learning of a naive optimist playing against a forward looking agent. I believe a generalization of such analysis might provide a foundation for self-confirming equilibria. Indeed, if both players care about consistency and generate beliefs about their opponent's strategy, experimentation may be limited. When players do not foresee the effects of their current choices of beliefs, then stickiness, or non-experimentation, may be even stronger. Hence, a formalization of these ideas would specify a learning process in which for sufficiently high γ of both players the play converges to self-confirming equilibria. The speed of convergence would be related to the difference between the actual parameter γ and the perceived $\tilde{\gamma}$.

On an applied level, a natural next step is a quantitative test of the model's predictions. In order to carry out such a test, one would need numerical estimates of the model's parameters, namely estimates of γ and $\tilde{\gamma}$. As indicated by some of the medical and psychological studies (see discussion in Section IV.D) these parameters are probably not uniform across individuals. Hence, a calibration exercise would entail a specification of distributions, rather than sheer means.

Appendix A - Proofs of Main Results

Lemma: *Expected utility at stage 1, if $q_2 < \tilde{\gamma}$, is:*

$$U = \begin{cases} (1 + \delta)\mu_1^B + \gamma[\mu_1(1 - \delta) + \delta - p] & \mu_1 \geq \frac{q_2(2 - \tilde{\gamma})}{2 + \tilde{\gamma} - 2q_2} \\ [\mu_1^B + \gamma(\mu_1 - p)] + \delta[q_2 - \frac{1}{2}\gamma(q_2 + \mu_1^B - 2) - \gamma\mu_1 + \gamma q_2 \mu_1^B] & \frac{1}{2} \leq \mu_1 < \frac{q_2(2 - \tilde{\gamma})}{2 + \tilde{\gamma} - 2q_2} \\ [1 - \mu_1^B + \gamma(\mu_1 - p)] + \delta[q_2 + \frac{1}{2}\gamma(q_2 + \mu_1^B - 1) + \gamma\mu_1 - \gamma q_2 \mu_1^B] & \frac{(2 + \tilde{\gamma})(1 - q_2)}{2 - \tilde{\gamma}(2q_2 - 1)} \leq \mu_1 < \frac{1}{2} \\ (1 + \delta)(1 - \mu_1^B + \gamma\mu_1) - \gamma p & \mu_1 < \frac{(2 + \tilde{\gamma})(1 - q_2)}{2 - \tilde{\gamma}(2q_2 - 1)} \end{cases}.$$

If $q_2 \geq \tilde{\gamma}$ then:

$$U = \begin{cases} (1 + \delta)\mu_1^B + \gamma[\mu_1(1 - \delta) + \delta - p] & \mu_1 \geq q_2 \\ [\mu_1^B + \gamma(\mu_1 - p)] + \delta[q_2 - \frac{1}{2}\gamma(q_2 + \mu_1^B - 2) - \gamma\mu_1 + \gamma q_2 \mu_1^B] & \frac{1}{2} \leq \mu_1 < q_2 \\ [1 - \mu_1^B + \gamma(\mu_1 - p)] + \delta[q_2 + \frac{1}{2}\gamma(q_2 + \mu_1^B - 1) + \gamma\mu_1 - \gamma q_2 \mu_1^B] & 1 - q_2 \leq \mu_1 < \frac{1}{2} \\ (1 + \delta)(1 - \mu_1^B + \gamma\mu_1) - \gamma p & \mu_1 < 1 - q_2 \end{cases}.$$

Proof: I use backwards induction in order to solve for equilibrium.

Looking at stage 2, assume first that $\mu_1 > \frac{1}{2}$.

If $q_2 \leq \mu_1$ then $\mu_2^B \geq \frac{1}{2}$ for all $s_2 \rightarrow \mu_2 = 1$ and $a_0 = a_1 = a_2 = L$.

If $q_2 > \mu_1$ then $\mu_2^B(s_2 = R) < \frac{1}{2} \rightarrow$ comparison between:

1. $\mu_2 = \frac{1}{2}$ and $a_2 = R$ yielding $\frac{(1 - \mu_1)q_2}{\mu_1(1 - q_2) + (1 - \mu_1)q_2} + \tilde{\gamma}(\frac{1}{2} - \mu_1)$.
2. $\mu_2 = 1$ and $a_2 = L$ yielding $\frac{\mu_1(1 - q_2)}{\mu_1(1 - q_2) + (1 - \mu_1)q_2} + \tilde{\gamma}(1 - \mu_1)$.

Hence,

$$\mu_2 = \begin{cases} \frac{1}{2} & \mu_1 < \frac{q_2(2 - \tilde{\gamma})}{2 + \tilde{\gamma} - 2q_2} \\ 1 & \mu_1 > \frac{q_2(2 - \tilde{\gamma})}{2 + \tilde{\gamma} - 2q_2} \end{cases}$$

(the agent is indifferent when $\mu_1 = \frac{q_2(2 - \tilde{\gamma})}{2 + \tilde{\gamma} - 2q_2}$).

Intuitively if μ_1 is sufficiently small the agent is more likely to value the signal R .

Assume now that $\mu_1 < \frac{1}{2}$.

If $q_2 \leq 1 - \mu_1$ then $\mu_2^B \leq \frac{1}{2}$ for all signals and $\mu_2 = 0$ and $a_1 = a_2 = R$.

If $q_2 > 1 - \mu_1$ then $\mu_2^B(s_2 = L) > \frac{1}{2} \rightarrow$ comparison between:

1. $\mu_2 = \frac{1}{2}$ and $a_2 = L$ that yields $\frac{\mu_1 q_2}{\mu_1 q_2 + (1 - \mu_1)(1 - q_2)} + \tilde{\gamma}(\mu_1 - \frac{1}{2})$.
2. $\mu_2 = 0$ and $a_2 = R$ that yields $\frac{(1 - \mu_1)(1 - q_2)}{\mu_1 q_2 + (1 - \mu_1)(1 - q_2)} + \tilde{\gamma}\mu_1$.

Optimization leads to:

$$\mu_2 = \begin{cases} \frac{1}{2} & \mu_1 > \frac{(2 + \tilde{\gamma})(1 - q_2)}{2 - \tilde{\gamma}(2q_2 - 1)} \\ 0 & \mu_1 < \frac{(2 + \tilde{\gamma})(1 - q_2)}{2 - \tilde{\gamma}(2q_2 - 1)} \end{cases}$$

(the agent is indifferent when $\mu_1 = \frac{(2+\tilde{\gamma})(1-q_2)}{2-\tilde{\gamma}(2q_2-1)}$).

Going back to stage 1, if $\mu_1 \geq \frac{1}{2}$ then if $q_2 < \tilde{\gamma}$, $\frac{q_2(2-\tilde{\gamma})}{2+\tilde{\gamma}-2q_2} < q_2$ and we have two cases to consider:
 If $\mu_1 < \frac{q_2(2-\tilde{\gamma})}{2+\tilde{\gamma}-2q_2}$, $\mu_2(s_2 = R) = \frac{1}{2}$ and $a_2(s_2 = R) = R$. The utility at stage 1 is:

$$[\mu_1^B + \gamma(\mu_1 - p)] + \delta[q_2 + \gamma - \frac{1}{2}\gamma(q_2 + \mu_1^B) - \gamma\mu_1 + \gamma q_2 \mu_1^B].$$

If $\frac{q_2(2-\tilde{\gamma})}{2+\tilde{\gamma}-2q_2} < \mu_1$ then $\mu_2 = 1$ and $a_2 = L$ no matter what s_2 is and the utility is:

$$[\mu_1^B + \gamma(\mu_1 - p)] + \delta[\mu_1^B + \gamma(1 - \mu_1)] = (1 + \delta)\mu_1^B + \gamma[\mu_1(1 - \delta) + \delta - p].$$

Summarizing,

$$U \quad | \quad \mu_1 \geq \frac{1}{2}, q_2 < \tilde{\gamma} (t = 2) = \begin{cases} [\mu_1^B + \gamma(\mu_1 - p)] + \delta[q_2 - \frac{1}{2}\gamma(q_2 + \mu_1^B - 2) - \gamma\mu_1 + \gamma q_2 \mu_1^B] & \mu_1 < \frac{q_2(2-\tilde{\gamma})}{2+\tilde{\gamma}-2q_2} \\ (1 + \delta)\mu_1^B + \gamma[\mu_1(1 - \delta) + \delta - p] & \mu_1 > \frac{q_2(2-\tilde{\gamma})}{2+\tilde{\gamma}-2q_2} \end{cases}$$

If $q_2 \geq \tilde{\gamma}$ then $\frac{q_2(2-\tilde{\gamma})}{2+\tilde{\gamma}-2q_2} \geq q_2$ and the expected utility at stage 1 is given by:

$$U \quad | \quad \mu_1 \geq \frac{1}{2}, q_2 \geq \tilde{\gamma} (t = 2) = \begin{cases} [\mu_1^B + \gamma(\mu_1 - p)] + \delta[q_2 - \frac{1}{2}\gamma(q_2 + \mu_1^B - 2) - \gamma\mu_1 + \gamma q_2 \mu_1^B] & \mu_1 < q_2 \\ (1 + \delta)\mu_1^B + \gamma[\mu_1(1 - \delta) + \delta - p] & \mu_1 > q_2 \end{cases}$$

Assume now that $\mu_1 < \frac{1}{2}$ then if $q_2 < \tilde{\gamma}$, $\frac{(2+\tilde{\gamma})(1-q_2)}{2-\tilde{\gamma}(2q_2-1)} > 1 - q_2$ and, as before, we have two cases to consider:

If $\mu_1 > \frac{(2+\tilde{\gamma})(1-q_2)}{2-\tilde{\gamma}(2q_2-1)}$ then the expected payoff at stage 1 is:

$$[1 - \mu_1^B + \gamma(\mu_1 - p)] + \delta[q_2 + \frac{1}{2}\gamma(q_2 + \mu_1^B - 1) + \gamma\mu_1 - \gamma q_2 \mu_1^B].$$

If $\mu_1 < \frac{(2+\tilde{\gamma})(1-q_2)}{2-\tilde{\gamma}(2q_2-1)}$ then $a_2 = R$ always and expected utility is:

$$[1 - \mu_1^B + \gamma(\mu_1 - p)] + \delta[1 - \mu_1^B + \gamma\mu_1] = (1 + \delta)(1 - \mu_1^B + \gamma\mu_1) - \gamma p.$$

Therefore,

$$\begin{aligned}
U \quad | \quad \mu_1 < \frac{1}{2}, q_2 < \tilde{\gamma} (t=2) &= \\
= \begin{cases} [1 - \mu_1^B + \gamma(\mu_1 - p)] + \delta[q_2 + \frac{1}{2}\gamma(q_2 + \mu_1^B - 1) + \gamma\mu_1 - \gamma q_2 \mu_1^B] & \mu_1 > \frac{(2+\tilde{\gamma})(1-q_2)}{2-\tilde{\gamma}(2q_2-1)} \\ (1 + \delta)(1 - \mu_1^B + \gamma\mu_1) - \gamma p & \mu_1 < \frac{(2+\tilde{\gamma})(1-q_2)}{2-\tilde{\gamma}(2q_2-1)} \end{cases} .
\end{aligned}$$

If $q_2 \geq \tilde{\gamma}$ then $\frac{(2+\tilde{\gamma})(1-q_2)}{2-\tilde{\gamma}(2q_2-1)} \leq 1 - q_2$ and the expected utility at stage 1 is given by:

$$\begin{aligned}
U \quad | \quad \mu_1 < \frac{1}{2}, q_2 \geq \tilde{\gamma} (t=2) &= \\
= \begin{cases} [1 - \mu_1^B + \gamma(\mu_1 - p)] + \delta[q_2 + \frac{1}{2}\gamma(q_2 + \mu_1^B - 1) + \gamma\mu_1 - \gamma q_2 \mu_1^B] & \mu_1 > 1 - q_2 \\ (1 + \delta)(1 - \mu_1^B + \gamma\mu_1) - \gamma p & \mu_1 < 1 - q_2 \end{cases} .
\end{aligned}$$

Hence, at stage 1, if $q_2 < \tilde{\gamma}$ then:

$$U = \begin{cases} (1 + \delta)\mu_1^B + \gamma[\mu_1(1 - \delta) + \delta - p] & \mu_1 \geq \frac{q_2(2-\tilde{\gamma})}{2+\tilde{\gamma}-2q_2} \\ [\mu_1^B + \gamma(\mu_1 - p)] + \delta[q_2 - \frac{1}{2}\gamma(q_2 + \mu_1^B - 2) - \gamma\mu_1 + \gamma q_2 \mu_1^B] & \frac{1}{2} \leq \mu_1 < \frac{q_2(2-\tilde{\gamma})}{2+\tilde{\gamma}-2q_2} \\ [1 - \mu_1^B + \gamma(\mu_1 - p)] + \delta[q_2 + \frac{1}{2}\gamma(q_2 + \mu_1^B - 1) + \gamma\mu_1 - \gamma q_2 \mu_1^B] & \frac{(2+\tilde{\gamma})(1-q_2)}{2-\tilde{\gamma}(2q_2-1)} \leq \mu_1 < \frac{1}{2} \\ (1 + \delta)(1 - \mu_1^B + \gamma\mu_1) - \gamma p & \mu_1 < \frac{(2+\tilde{\gamma})(1-q_2)}{2-\tilde{\gamma}(2q_2-1)} \end{cases} .$$

If $q_2 \geq \tilde{\gamma}$ then:

$$U = \begin{cases} (1 + \delta)\mu_1^B + \gamma[\mu_1(1 - \delta) + \delta - p] & \mu_1 \geq q_2 \\ [\mu_1^B + \gamma(\mu_1 - p)] + \delta[q_2 - \frac{1}{2}\gamma(q_2 + \mu_1^B - 2) - \gamma\mu_1 + \gamma q_2 \mu_1^B] & \frac{1}{2} \leq \mu_1 < q_2 \\ [1 - \mu_1^B + \gamma(\mu_1 - p)] + \delta[q_2 + \frac{1}{2}\gamma(q_2 + \mu_1^B - 1) + \gamma\mu_1 - \gamma q_2 \mu_1^B] & 1 - q_2 \leq \mu_1 < \frac{1}{2} \\ (1 + \delta)(1 - \mu_1^B + \gamma\mu_1) - \gamma p & \mu_1 < 1 - q_2 \end{cases} \blacksquare$$

Proof of Proposition 1: For sufficiently high γ ,

$$(1 + \delta) + \gamma[(1 - \delta) + \delta - p] > [1 + \gamma(\frac{q_2(2 - \tilde{\gamma})}{2 + \tilde{\gamma} - 2q_2} - p)] + \delta[q_2 - \frac{1}{2}\gamma(q_2 + 1 - 2) - \gamma\frac{q_2(2 - \tilde{\gamma})}{2 + \tilde{\gamma} - 2q_2} + \gamma q_2],$$

$$(1 + \delta) + \gamma[(1 - \delta) + \delta - p] > [1 + \gamma(q_2 - p)] + \delta[q_2 - \frac{1}{2}\gamma(q_2 + 1 - 2) - \gamma q_2 + \gamma q_2],$$

$$(1 + \delta)(1 - 0 + \gamma\frac{(2 + \tilde{\gamma})(1 - q_2)}{2 - \tilde{\gamma}(2q_2 - 1)}) - \gamma p > [1 - 0 + \gamma(\frac{1}{2} - p)] + \delta[q_2 + \frac{1}{2}\gamma(q_2 + 0 - 1) + \gamma\frac{1}{2} - \gamma q_2 0],$$

and

$$(1 + \delta)(1 - 0 + \gamma(1 - q_2)) - \gamma p > [1 - 0 + \gamma(\frac{1}{2} - p)] + \delta[q_2 + \frac{1}{2}\gamma(q_2 + 0 - 1) + \gamma\frac{1}{2} - \gamma q_2 0].$$

Using the Lemma, this implies that when $q_1 = 1$, the agent will choose the most extreme beliefs -

$$\mu_1(s_1 = L) = 1, \quad \mu_1(s_1 = R) = \begin{cases} \frac{(2+\tilde{\gamma})(1-q_2)}{2-\tilde{\gamma}(2q_2-1)} & q_2 \leq \tilde{\gamma} \\ 1 - q_2 & q_2 > \tilde{\gamma} \end{cases}.$$

From continuity, these equalities will hold for a range of q_1 's, namely all $q_1 > \tilde{q}_1(\gamma)$ where $\tilde{q}_1(\gamma) \in (\frac{1}{2}, 1)$.

Define q'_1 to satisfy $1 - \frac{p(1-q'_1)}{p(1-q'_1)+(1-p)q'_1} = 1 - q_2$ so that when $q_1 = q'_1$, $\mu_1^B(s_1 = R) = q_2$. Note that $q'_1 < 1$ whenever $p < 1$.

Choose $q_1^*(\gamma) = \max\{\tilde{q}_1(\gamma), q'_1\}$. $q_1^*(\gamma) < 1$ whenever $p < 1$. For all $q_1 > q_1^*(\gamma)$, $1 - q_2$, $\frac{(2+\tilde{\gamma})(1-q_2)}{2-\tilde{\gamma}(2q_2-1)} > 1 - \mu_1^B(s_1 = R)$ and so indeed the agent is overconfident when $s_1 = L$ and underconfident when $s_1 = R$. ■

Proof of Proposition 2: Using the notation of the last proposition, when $q_1 > q_1(\gamma)$, $\mu_1(s_1 = L) > q_2$ and from the proof of the Lemma, $\mu_2 = 1$ no matter what s_2 is. Since $q_2 > \tilde{\gamma}$, $\mu_1(s_1 = R) = 1 - q_2$. Hence, from the proof of the Lemma, $\mu_2 = 0$ no matter what s_2 is. Therefore,

$$a_1(s_1 = L) = a_2(s_1 = L, s_2 = L) = a_2(s_1 = L, s_2 = R) = L;$$

$$a_1(s_1 = R) = a_2(s_1 = R, s_2 = L) = a_2(s_1 = R, s_2 = R) = R. \blacksquare$$

Proof of Proposition 3.1 (ex-ante utility non-monotonicity w.r.t. q_1): We look at the case of $\delta = 1, \gamma = 1, \tilde{\gamma} = 0$. Choose $q_2 > p$, q_1^L s.t. $\mu_1(s_1 = L) = q_2$ and $\mu_1(s_1 = R) = \frac{1}{2}$ and then choose $q_1^H (> q_1^L)$ s.t. $\mu_1(s_1 = L) = 1$ and $\mu_1(s_1 = R) = 0$.

The ex-ante expected utility when $q_1 = q_1^L$ is:

$$\begin{aligned} & pq_1^L [2 + (q_2 - p) + (1 - q_2)] + p(1 - q_1^L) [q_2 + (\frac{1}{2} - p) + (1 - \frac{1}{2})q_2] + \\ & + (1 - p)q_1^L [1 + q_2 + (\frac{1}{2} - p) + (1 - \frac{1}{2})(1 - q_2)] + (1 - p)(1 - q_1^L) [q_2 - p + 1 - q_2] \\ & = pq_1^L (3 - p) + p(1 - q_1^L) (\frac{3}{2}q_2 + \frac{1}{2} - p) + (1 - p)q_1^L (2 + \frac{1}{2}q_2 - p) + (1 - p)(1 - q_1^L)(1 - p). \end{aligned}$$

The ex-ante expected utility when $q_1 = q_1^H$ is:

$$\begin{aligned} & pq_1^H[2 + (1 - p)] + p(1 - q_1^H)[0 - p] + (1 - p)q_1^H[2 + (0 - p)] + (1 - p)(1 - q_1^H)[1 - p] = \\ & = pq_1^H(3 - p) + p(1 - q_1^H)(-p) + (1 - p)q_1^H(2 - p) + (1 - p)(1 - q_1^H)(1 - p). \end{aligned}$$

subtracting the latter from the former we get:

$$\begin{aligned} & (q_1^H - q_1^L)[(1 - p)^2 - p(3 - p) - p^2 - (1 - p)(1 - 2p)] + \\ & + p(1 - q_1^L)\left(\frac{3}{2}q_2 + \frac{1}{2}\right) + (1 - p)q_1^L\frac{1}{2}q_2 \\ & = 2p(p - 1)(q_1^H - q_1^L) + \frac{1}{2} + \frac{3}{2}pq_2 + \frac{1}{2}q_1^Lq_2 - 2pq_1^Lq_2 - \frac{1}{2}pq_1^L \end{aligned}$$

From the assumptions on the parameters,

$$2p(1 - p)(q_1^H - q_1^L) < \frac{1}{4},$$

$$\frac{3}{2}pq_2 - 2pq_1^Lq_2 = \frac{3}{2}pq_2(1 - q_1^L) - \frac{1}{2}pq_1^Lq_2 > -\frac{1}{2}pq_1^Lq_2 > -\frac{1}{4}, \text{ and}$$

$$\frac{1}{2}q_1^Lq_2 - \frac{1}{2}pq_1^L = \frac{1}{2}q_1^L(q_2 - p) > 0.$$

Thus, the above difference is positive and q_1^L leads to higher ex-ante expected utility than q_1^H . The claim of the Proposition follows from continuity. ■

Proof of Proposition 3.2 (ex-ante utility non-monotonicity w.r.t. q_2): I look at the case of $\delta = 1$, $\gamma = 1$, $\tilde{\gamma} = 0$ and compare the ex-ante expected utility when $q_2 = \frac{1}{2}$ and $q_2 = 1 - \varepsilon$ where ε is small.

For $q_2 = \frac{1}{2}$, when choosing a belief, the agent believed that she will not alter her belief in the next period. Since no additional information will be acquired, the instrumental term of the utility is the same for both periods, and the belief is predicted to equal zero in the second period. Hence, it is as if the agent is solving a one period game with $\gamma = \frac{1}{2}$. Using our analysis for the Proof of the Lemma above, since $q_1 < \frac{3q_1}{5-4q_1}$, we get that when $q_1 > p$:

$$\mu_1(s_1 = L) = 1 \quad \mu_1(s_1 = R) = \begin{cases} \frac{1}{2} & p < \frac{3q_1}{5-4q_1} \\ 1 & p > \frac{3q_1}{5-4q_1} \end{cases} = \frac{1}{2}.$$

Expected ex-ante utility is then:

$$pq_1[2 + (1 - p)] + p(1 - q_1)[\frac{1}{2} - p] + (1 - p)q_1[2 + \frac{1}{2} - p] + (1 - p)(1 - q_1)(1 - p).$$

For $q_2 = 1 - \varepsilon$, $\varepsilon \ll 1$, we can use the Lemma in order to determine the choice of μ_1 . The agent expects total utility to equal:

$$U = \begin{cases} 2\mu_1^B + \gamma(1 - p) & \mu_1 \geq 1 - \varepsilon \\ [\mu_1^B + \gamma(\mu_1 - p)] + [1 + \frac{1}{2}\gamma(1 - \mu_1^B) + \gamma(\mu_1^B - \mu_1)] & \frac{1}{2} \leq \mu_1 < 1 - \varepsilon \\ [1 - \mu_1^B + \gamma(\mu_1 - p)] + [1 + \frac{1}{2}\gamma\mu_1^B - \gamma(\mu_1^B - \mu_1)] & \varepsilon \leq \mu_1 < \frac{1}{2} \\ 2(1 - \mu_1^B + \gamma\mu_1) - \gamma p & \mu_1 < \varepsilon \end{cases}$$

For sufficiently high q_1 we get:

$$\mu_1 = \begin{cases} 1 & s_1 = L \\ \varepsilon & s_1 = R \end{cases}.$$

At stage 2, since $q_2 = 1 - \varepsilon (< 1)$,

$$\mu_2 = \begin{cases} \mu_1 = 1 & s_1 = L \\ 0 & s_1 = R \end{cases}$$

and expected utility is therefore:

$$pq_1[2 + (1 - p)] + p(1 - q_1)(2\varepsilon - p) + (1 - p)q_1(2 + 2\varepsilon - p) + (1 - p)(1 - q_1)(1 - p).$$

The difference between the ex-ante expected utility with $q_2 = \frac{1}{2}$ and ex-ante expected utility with $q_2 = 1 - \varepsilon$, $\varepsilon \ll 1$ is then approximately (when letting $\varepsilon \rightarrow 0$)

$$\begin{aligned} p(1 - q_1)(\frac{1}{2} - p + p) + (1 - p)q_1(2 + \frac{1}{2} - p - 2 + p) &= \\ = \frac{1}{2}p(1 - q_1) + \frac{1}{2}(1 - p)q_1 &> 0. \end{aligned}$$

So, indeed, the ex-ante expected utility is higher for $q_2 = \frac{1}{2}$ than for $q_2 = 1 - \varepsilon > \frac{1}{2}$. The claim of the proposition follows from continuity. ■

Proof of Proposition 4.1 (perceived utility monotonicity w.r.t. q_1): The proof is essentially identical to the one used to show monotonicity of ex-ante utility when agents are forward look-

ing (see Appendix B that follows). Let $r_1 > q_1 \geq \frac{1}{2}$. Let $y \in [\frac{1}{2}, 1]$ be s.t.

$$r_1 y + (1 - r_1)(1 - y) = q_1.$$

Consider the following strategy in the game with signals of accuracies r_1, q_2, \dots, q_T :

After receiving the signal s_1 create a signal $\bar{s}_1 \in \{L, R\}$ according to the probabilities:

$$\Pr(\bar{s}_1 = s_1) = y \quad \Pr(\bar{s}_1 \neq s_1) = 1 - y.$$

From the choice of y ,

$$\Pr(\bar{s}_1 = \theta) = q_1 \quad \Pr(\bar{s}_1 \neq \theta) = 1 - q_1.$$

Thus, \bar{s}_1 is a signal of accuracy q_1 . The agent can then use the optimal strategy of the game with signal accuracies q_1 and can hence guarantee herself a perceived utility level corresponding to accuracies (q_1, q_2) when the accuracies are (r_1, q_2) . ■

Proof of Proposition 4.2 (perceived utility non-monotonicity w.r.t. q_2): Suppose that $\delta = 1, \gamma = 1$, and $\tilde{\gamma}, q_1, q_2^L, q_2^H, q_2^L < \tilde{\gamma} < q_2^H$ are such that

1.

$$\frac{1}{2}q_2^L - \frac{3}{2}q_1 + q_1q_2^H > 0,$$

2. When $q_2 = q_2^L$,

$$\mu_1 = \begin{cases} \frac{q_2(2-\tilde{\gamma})}{2+\tilde{\gamma}-2q_2} & s_1 = L \\ \frac{1}{2} & s_1 = R \end{cases} \quad a_2 = \begin{cases} L & s_2 = L \\ R & s_2 = R \end{cases},$$

3. When $q_2 = q_2^H$,

$$\mu_1 = \begin{cases} q_2 & s_1 = L \\ \frac{1}{2} & s_1 = R \end{cases} \quad a_2 = \begin{cases} L & s_1 = L \text{ or } s_2 = L \\ R & s_1 = s_2 = R \end{cases}.$$

When $q_2 = q_2^L$, expected perceived utility is:

$$p\{q_1[1 + \frac{q_2^L(2-\tilde{\gamma})}{2+\tilde{\gamma}-2q_2^L} - p + q_2^L(2 - \frac{q_2^L(2-\tilde{\gamma})}{2+\tilde{\gamma}-2q_2^L}) + (1 - q_2^L)(\frac{1}{2} - \frac{q_2^L(2-\tilde{\gamma})}{2+\tilde{\gamma}-2q_2^L})\} +$$

$$\begin{aligned}
& +(1 - q_1)\left[\frac{1}{2} - p + q_2^L\left(2 - \frac{1}{2}\right) + (1 - q_2^L)\left(0 + \frac{1}{2}\right)\right] + \\
& +(1 - p)\left\{q_1\left[1 + \frac{1}{2} - p + q_2^L\left(1 + \frac{1}{2}\right) + (1 - q_2^L)\frac{1}{2}\right] + \right. \\
& \left. + (1 - q_1)\left[\frac{q_2^L(2 - \tilde{\gamma})}{2 + \tilde{\gamma} - 2q_2^L} - p + q_2^L\left(1 + \frac{1}{2} - \frac{q_2^L(2 - \tilde{\gamma})}{2 + \tilde{\gamma} - 2q_2^L}\right) + (1 - q_2^L)\left(1 - \frac{q_2^L(2 - \tilde{\gamma})}{2 + \tilde{\gamma} - 2q_2^L}\right)\right]\right\} = \\
& = 1 - p + q_1 + \frac{1}{2}q_2^L - \frac{1}{2}pq_1 + \frac{1}{2}pq_2^L + \frac{1}{2}q_1q_2^L.
\end{aligned}$$

When $q_2 = q_2^H$, expected perceived utility is:

$$\begin{aligned}
& p\{q_1[2 + q_2 - p + 1 - q_2^H] + (1 - q_1)[0 + \frac{1}{2} - p + q_2^H(1 + \frac{1}{2}) + (1 - q_2^H)(0 + \frac{1}{2})]\} + \\
& +(1 - p)\{q_1[1 + \frac{1}{2} - p + q_2^H(1 + \frac{1}{2}) + (1 - q_2^H)(0 + \frac{1}{2})] + \\
& \quad + (1 - q_1)[0 + q_2 - p + 1 - q_2^H]\} = \\
& = 1 - p + q_1 + pq_1 + q_1q_2^H(1 - p)
\end{aligned}$$

Subtracting the latter from the former we get:

$$\begin{aligned}
& \frac{1}{2}q_2^L - \frac{1}{2}pq_1 + \frac{1}{2}pq_2^L + \frac{1}{2}q_1q_2^L - pq_1 - q_1q_2^H(1 - p) = \\
& = p\left(\frac{1}{2}q_2^L - \frac{3}{2}q_1 + q_1q_2^H\right) + \frac{1}{2}q_2^L(1 + q_1) - q_1q_2^H
\end{aligned}$$

since $\frac{1}{2}q_2^L - \frac{3}{2}q_1 + q_1q_2^H > 0$, we get that the last expression is greater or equal than the expression achieved after plugging in $p = \frac{1}{2}$. That is, the difference is greater or equal than:

$$\begin{aligned}
& \frac{1}{2}\left(\frac{1}{2}q_2^L - \frac{3}{2}q_1 + q_1q_2^H\right) + \frac{1}{2}q_2^L(1 + q_1) - q_1q_2^H = \\
& = \frac{1}{2}q_1(q_2^L - q_2^H) + \frac{3}{4}(q_2^L - q_1) \geq -\frac{1}{4} + \frac{3}{8} > 0
\end{aligned}$$

where the inequality before last holds since $q_1 \leq 1$, $q_2^L - q_2^H \geq -\frac{1}{2}$, and $q_2^L - q_1 \leq \frac{1}{2}$.

The proposition's claim then follows from continuity. ■

Proof of Proposition 5: We prove the claim for $\gamma = \tilde{\gamma} = 1, \delta = 1$. The proposition then follows from continuity. For the one stage game, equilibrium beliefs are

$$\mu_1(s_1 = L) = 1 \quad \mu_1(s_1 = R) = \begin{cases} 1 & q_1 \leq \frac{3p}{1+2p} \\ \frac{1}{2} & q_1 > \frac{3p}{1+2p} \end{cases} .$$

The ex-ante expected utility of the one stage game is given by:

$$U_1^* = \begin{cases} 1 & q_1 \leq \frac{3p}{1+2p} \\ \frac{1}{2} - p + q_1 + \frac{1}{2}pq_1 & q_1 > \frac{3p}{1+2p} \end{cases} .$$

For any $q_2 < \frac{3}{4}$,

$$\mu_2 = \begin{cases} 1 & \mu_1 \geq \frac{1}{2} \\ 0 & \mu_1 < \frac{1}{2} \end{cases} .$$

Thus, the ex-ante expected utility is:

$$U_2^* \leq [pq_1 + (1-p)q_1] + \frac{1}{2}(1-p) = \frac{1}{2}(1-p) + q_1 .$$

Since $q_1 < \frac{3p}{1+2p}, \frac{1+p}{2}$, we get that $U_2^* < U_1^*$ and the proposition's claim follows from continuity. ■

Appendix B - Forward Looking Agents

The case of forward looking agents, i.e. agents for whom $\gamma = \tilde{\gamma}$ turns out conceptually different from the case of $\tilde{\gamma} < \gamma$. These agents are time consistent and therefore standard results of information economics, such as preferring more information to less, preferring more accurate signals over less accurate ones, etc. hold for them. In this appendix I deal with the special results relating to forward looking agents. The phenomena characteristic of forward looking agents, that I present here, are not specific to the utility functional form I pick, or to the length of the process. I will hence phrase them in the context of a general functional form as:

$$U_t = \sum_{i=t}^T \delta^{i-t} u(a_i) + \sum_{j=t-1}^T \delta^{i-t} v(\mu_{t-1}, \mu_t) \quad \mu_{-1} \equiv \mu_0 .$$

Ex-ante expected utility is the expected utility before any signal is realized.

Denote by $U(q_1, q_2, \dots, q_T)$ the ex-ante expected utility with signals of accuracies q_1, q_2, \dots, q_T . For forward looking agents, the ex-ante expected utility is the utility the agent of time zero expects to get, before observing signal s_1 . For such agents, there is free disposal of information, in the sense

that information can be ignored or distorted (and made less accurate). Hence, ex-ante utility is non-decreasing in the in each of the signals' accuracies. This is captured in the following proposition:

Proposition 8 (general monotonicity of ex-ante expected utility): $U(q_1, q_2, \dots, q_T)$ is non-decreasing in q_i for all i over $[\frac{1}{2}, 1]^T$. In particular, if U is differentiable in all its arguments, then $\frac{\partial U}{\partial q_i} \geq 0$ for all i .

Proof: Let $r_1 > q_1 \geq \frac{1}{2}$. Let $y \in [\frac{1}{2}, 1]$ be s.t.

$$r_1 y + (1 - r_1)(1 - y) = q_1.$$

Consider the following strategy in the game with signals of accuracies r_1, q_2, \dots, q_T :

After receiving the signal s_1 create a signal $\bar{s}_1 \in \{L, R\}$ according to the probabilities:

$$\Pr(\bar{s}_1 = s_1) = y \quad \Pr(\bar{s}_1 \neq s_1) = 1 - y.$$

From the choice of y ,

$$\Pr(\bar{s}_1 = \theta) = q_1 \quad \Pr(\bar{s}_1 \neq \theta) = 1 - q_1.$$

Thus, \bar{s}_1 is a signal of accuracy q_1 . The agent then uses the optimal strategy of the game with signal accuracies q_1, q_2, \dots, q_T using the signals $\bar{s}_1, s_2, \dots, s_T$ thus guaranteeing an expected (ex-ante) utility of $U(q_1, q_2, \dots, q_T)$. Hence,

$$U(r_1, q_2, \dots, q_T) \geq U(q_1, q_2, \dots, q_T)$$

as the proposition claims. ■

Proposition 9 (convexity of ex-ante expected utility): $U(q_1, q_2, \dots, q_T)$ is convex in each of q_1, q_2, \dots, q_T .

Proof: Let $q_i = \alpha q_i^1 + (1 - \alpha)q_i^2$. Consider the game in which before the process begins, accuracies q_i^1 and q_i^2 are randomized with probabilities α and $1 - \alpha$ respectively and the agent is told which accuracy was realized. The ex-ante expected equilibrium utility of this game is $\alpha U(q_1, \dots, q_{i-1}, q_i^1, q_{i+1}, \dots, q_T) + (1 - \alpha)U(q_1, \dots, q_{i-1}, q_i^2, q_{i+1}, \dots, q_T)$. If the agent ignores the additional information concerning which accuracy has been chosen, she can guarantee $U(q_1, \dots, q_{i-1}, q_i, q_{i+1}, \dots, q_T)$. Thus, the definition of equilibrium strategy requires:

$$\begin{aligned} \alpha U(q_1, \dots, q_{i-1}, q_i^1, q_{i+1}, \dots, q_T) + (1 - \alpha)U(q_1, \dots, q_{i-1}, q_i^2, q_{i+1}, \dots, q_T) &\geq \\ &\geq U(q_1, \dots, q_{i-1}, q_i, q_{i+1}, \dots, q_T) \end{aligned}$$

which implies that U is convex in its arguments. ■

Corollary: $U(q_1, q_2, \dots, q_T)$ is continuous on $(\frac{1}{2}, 1)^T$.

References

- Akerlof, G. A. and Dickens, W. T. [1982], "The Economic Consequences of Cognitive Dissonance," *The American Economic Review*, 72(3), pp. 307-319.
- Akerlof, G. A. and Kranton, R. [2000], "Economics and Identity," *The Quarterly Journal of Economics*.
- Ansolabehere, S. and Iyengar, S. [1994], "Riding the Wave and Claiming Ownership Over Issues: The Joint Effects of Advertising and News Coverage in Campaigns," *Public Opinion Quarterly*, 58(3), pp. 335-357.
- Arkes, H. R. and Blumer, C. [1985], "The Psychology of Sunk Costs," *Organizational Behavior and Human Decision Processes*, 35, pp. 124-140.
- Aronson, E. [1969], "The Theory of Cognitive Dissonance: A Current Perspective." In L. Berkowitz (Ed.), *Advances in Experimental Social Psychology* (volume 4), New York: Academic.
- Baumeister, R. F. [1998], "The Self," in Gilbert, D. T., Fiske, S. T., and Lindzey, G. (Eds.), *The handbook of social psychology* (4th edition), New York: McGraw Hill.
- Benabou, R. and Tirole, J. [2002], "Self-Confidence and Personal Motivation," forthcoming in *The Quarterly Journal of Economics*.
- Blackwell, D. [1950], "Equivalent Comparisons of Experiments," *Annals of Mathematical Statistics*, 21, 265-272.
- Bodner, R. and Prelec, D. [1997], "The Diagnostic Value of Actions in a Self-signaling Model," mimeo.
- Brehm, J. W. [1956], "Postdecision Changes in the Desirability of Alternatives," *Journal of Abnormal Social Psychology*, 52, pp. 384-389.
- Carillo, J. and Mariotti, T. [1997], "Wishful Thinking and Strategic Ignorance," mimeo.
- Darley, J. M. and Gross, P. H. [1983], "A Hypothesis-Confirming Bias in Labeling Effects," *Journal of Personality and Social Psychology*, pp. 20-33.
- Ehrlich, D., Guttman, I., Schonbach, P., and Mills, J. [1957], "Post-decision Exposure to Relevant Information," *Journal of Abnormal and Social Psychology*, 54, pp. 98-102.
- Elkin, R. H. and Leippe, M. R. [1986], "Physiological Arousal, Dissonance, and Attitude Change: Evidence for a Dissonance-Arousal Link and a "Don't Remind Me" Effect," *Journal of Personality and Social Psychology*, 51, pp. 55-65.
- Elliot, A. J. and Devine, P. G. [1994], "On the Motivational Nature of Cognitive Dissonance: Dissonance as Psychological Discomfort," *Journal of Personality and Social Psychology*, 67(3), pp. 382-394.
- Festinger, L. [1957], *A Theory of Cognitive Dissonance*. Stanford, CA: Stanford.

Frank, M. G. and Gilovich, T. [1988], "The Dark Side of Self and Social Perception: Black Uniforms and Aggression in Professional Sports," *Journal of Personality and Social Psychology*, 54, pp. 74-85.

Freedman, J. L. and Fraser, S. C. [1966], "Compliance Without Pressure: The Foot-In-The-Door Technique," *Journal of Personality and Social Psychology*, 4, pp. 195-202.

Fromkin, V. [2000], *Linguistics: An Introduction to Linguistic Theory*, Oxford, UK: Blackwell.

Geanakoplos, J., Pearce, D., and Stacchetti, E. [1989], "Psychological games and sequential rationality", *Games and Economic Behavior*, 1, 60-79.

Gilbert, D. T., Pinel, E. C., Wilson, T. D., Blumberg, S. J., and Wheatley, T. [1998], "Immune neglect: A source of durability bias in affective forecasting," *Journal of Personality and Social Psychology*, 75, pp. 617-638.

Gilovich, T. [1991], *How We Know What Isn't So: The Fallibility of Human Reason in Everyday Life*, Free Press.

Glass, D. C. [1964], "Changes in Liking as a Means of Reducing Cognitive Discrepancies between Self-Esteem and Aggression," *Journal of Personality*, 32, pp. 531-549.

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

Knox, R. E. and Inkster, J. A. [1968], "Postdecision Dissonance at Post Time," *Journal of Personality and Social Psychology*, 1(8), pp. 319-323.

Koszegi, B. [1999], "Self-image and Economic Behavior," MIT dissertation.

Krantz, D. S., Baum, A., and Weiderman, M. V. [1980], "Assessment of Preferences for Self-treatment and Information in Health Care," *Journal of Personality and Social Psychology*, 39, pp. 977-990.

Kunda, Z. [1990], "The Case of Motivated Reasoning," *Psychological Bulletin*, 108, 3, pp. 480-498.

Lazarsfeld, P. F., Berelson, B., and Gaudet, H. [1948], *The People's Choice*, Second Edition, New York, Columbia University Press.

Loewenstein, G. and Schkade, D. [1999], "Wouldn't it be nice? Predicting future feelings," In D. Kahneman, E. Diener, and N. Schwartz. (Eds.), *Well-being: The foundations of hedonic psychology*, Russell Sage Foundation.

Loewenstein, G., O'Donoghue, T. and Rabin, M. [2000], "Projection bias in predicting future utility," mimeo.

Loomes, G. and Sugden, R. [1982], "Regret Theory: An Alternative Theory of Rational Choice Under Uncertainty," *Economic Journal*, 92, pp. 804-824.

Lord, C., Lepper, M. R., and Ross, L. [1979], "Biases Assimilation and Attitude Polarization:

The Effects of Prior Theories on Subsequently Considered Evidence,” *Journal of Personality and Social Psychology*, 37(11), pp. 2098-2110.

Morgan, J., Roufeil, L., Kaushik, S., and Bassett, M. [1998], “Influence of Coping Style and Precolonoscopy Information on Pain and Anxiety of Colonoscopy,” *Gastrointestinal Endoscopy*, 48(2), pp. 119-127.

Nagler, M. G. [1993], “Rather bait than switch - Deceptive advertising with bounded consumer rationality,” *Journal of Public Economics*, 51, pp. 359-378.

Nisbett, R. E. and Ross, L. [1991], *Person and The Situation - Perspectives of Social Psychology*, Philadelphia: Temple University Press.

Oskamp, S. [1965], “Overconfidence in Case-Study Judgements,” *The Journal of Consulting Psychology*, 29, pp. 261-265.

Pelham, B. W. and Swann, W. [1989], “From self-conceptions to self-worth: The sources and structure of self-esteem,” *Journal of Personality and Social Psychology*, 57, pp. 672-680.

Pliner, P., Hart, H., Kohl, J., and Saari, D. [1974], “Compliance without Pressure: Some Further Data on the Foot-In-The-Door Technique,” *Journal of Experimental Social Psychology*, 10, pp. 17-22.

Rabin, M. and Schrag, J. L. [1999], “First Impressions Matter: A Model of Confirmatory Bias,” *The Quarterly Journal of Economics*, pp. 37-82.

Schaacter, D. L. [1996], *Searching for memory: The brain, the mind, and the past*, New York: Basic Books.

Sears, D. O. and Friedman, J. L. [1967], “Selective Exposure to Information: A Critical Review,” *Public Opinion Quarterly*, 31(2), pp. 194-213.

Shleifer, A. [2000], *Inefficient Markets*, Clarendon Lectures, Oxford University Press, 2000.

Staw, B. M. [1976], “Knee-deep in the Big Muddy: A Study of Escalation Commitment to a Chosen Course of Action,” *Organizational Behavior and Human Performance*, 16, pp. 27-44.

Swann, W. B. [1985], “The Self as an Architect of Social Reality,” in Schlenker, B. R. (Ed.), *The Self and Social Life*, New York: McGraw Hill.

Swann, W. B., De La Ronde, C., and Hixon, J. G. [1994], “Authenticity and Positivity Strivings in Marriage and Courtship,” *Journal of Personality and Social Psychology*, 66, pp. 857-869.

Tesser, A. [1988], “Toward a Self-Evaluation Maintenance Model of Social Behavior,” in Berkowitz, L. (Ed.), *Advances in Experimental Social Psychology*, San Diego, CA: Academic Press.

Tetlock, P. E., Skitka, L., and Boettger, R. [1989], "Social and Cognitive Strategies for Coping with Accountability - Conformity, Complexity, and Bolstering," *Journal of Personality and Social Psychology*, 57(4), pp. 632-640.

Yariv, L. [2001], "Believe and Let Believe: Axiomatic Foundations for Belief Dependent Utility Functionals," Cowles Foundation Discussion Paper Number 1344.