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

LEEAT YARIV*†

Current Version: January 4, 2005

ABSTRACT. Many observations from psychology, political science, and organizational behavior indicate 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. The current paper explores the implications of a simple model incorporating these effects into an agent's utility function. The model allows a characterization of when: 1. agents become under- and over-confident, 2. agents prefer less accurate signals, i.e., they are willing to pay in order to forgo information, and 3. agents exhibit excess stickiness or excess volatility in action choices.

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

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

*Department of Economics, University of California at Los Angeles and Division of Humanities and Social Sciences, Caltech. E-mail: lyariv@econ.ucla.edu

†I thank Drew Fudenberg, David Laibson, and Eric Maskin for many conversations and valuable suggestions. Bart Lipman and Muriel Niederle provided numerous helpful comments on a previous draft. I have also benefitted from talking with Alberto Alesina, Itzhak Gilboa, Ed Glaeser, Orit Kedar, David Levine, Ted Miguel, Paul Milgrom, Markus Mobius, Al Roth, Michael Schwarz, and Steve Tadelis. I am grateful to the Alfred P. Sloan foundation and the Eliot committee for generous financial support.

"I'll see it when I believe it."

-slip of tongue by experimental social psychologist Thane Pittman.

1. Introduction

There is abundant multi-disciplinary evidence indicating people's penchant for appearing consistent, both to themselves and to others. For instance, psychologists have identified agents' proclivity to exhibit *cognitive dissonance,* a tendency to change beliefs about the relative agreeableness of actions that have been taken, as well as *confirmatory bias,* the inclination to interpret new evidence in ways that confirm current beliefs; Political scientists often explore the implications of *selective exposure*, voters' propensity to seek out information (or interpret it in a way) that coincides with their ideological stands; Organizational behaviorists study the notion of *escalation of commitment*, which leads managers to persist with strategies that are revealed to be suboptimal.

The plethora of data accumulated in these fields ultimately suggests that people behave as if they have a taste for: 1. consistency between the action taken and the belief held at each point in a decision process and 2. consistency between beliefs in different stages of a decision process.

The current paper aims at incorporating these two effects in a simple dynamic model of decision-making. Specifically, I explore the implications generated by a model in which utility functions depend directly on beliefs, and beliefs are, in turn, objects of choice. Decision theoretical results can then be applied to agents' beliefs, in addition to agents' physical actions. The suggested framework helps explain an assortment of experimental observations related to information economics. In particular, I characterize when agents become overconfident or underconfident, identify cases in which agents are likely to persist with the action they chose initially, and show that agents may sometimes prefer less accurate signals over more accurate ones.

I consider a separable utility function that is comprised of two terms. The first is an *instrumental utility*, which coincides with the utility considered in standard economic analysis. The second term is a *belief utility*, corresponding to the direct utility from the sequence of beliefs adopted by the agent. This form of utility functions captures the trade-off between making optimal choices according to standard analysis (maximizing the instrumental utility) and being consistent (maximizing the belief utility).

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.[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 role in the outcomes of the decision process. Namely, it provides an avenue for time inconsistency that is not due to changing discount factors. I parametrize introspection with one (continuous) variable that denotes the weight the agent thinks she will place on the belief utility in the future. In particular, this allows the analysis to include cases in which agents perceive perfectly their future taste for consistency, or are utterly oblivious to it.

When agents have a taste for consistency and are sufficiently 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 2 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, discount the future sufficiently heavily, or have a sufficiently inaccurate second signal, then they are overconfident when the initial signal supports their prior belief and underconfident when their initial signal opposes it.

Biases in beliefs have a direct effect on the actions taken. Intuitively, consider an agent

[1] Everyday language suggests the resemblance of beliefs and standard actions. E.g., we talk about "holding beliefs," "giving up beliefs," "changing beliefs," etc.

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 her future updating process by much. However, once a later period arrives, the agent does care about consistency, does not update her beliefs sufficiently, and thereby 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 5 gives a formal characterization of situations in which such excess persistence arises, as well as identifies circumstances in which excess volatility occurs.

Looking at expected utility levels and the demand for information, I derive general results taking two points of view. The first is paternalistic. If an outsider can fully predict the agent's behavior, I analyze the kind of signals she would choose in order to maximize the agent's expected utility at the outset of the game. 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). 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, and excess persistence or volatility in action choices are phenomena that have been identified in economic behavior (see Section 2). The taste for less accurate information, as captured by Proposition 4 corresponds to observations that have hardly been empirically tested for in economic environments. Hence, the current framework both organizes several well-known empirical observations, as well as provides new predictions that can serve as a natural testing ground.

The structure of the paper is as follows. Section 2 provides some motivating evidence from psychology, political science, and organizational behavior for agents' taste for consistency, as well as reviews the related literature in economics. Section 3 contains the general setup of the analyzed dynamic decision problem. Section 4 presents the main results for the restricted setup. Sections 4.A and 4.B examine the choice of beliefs and actions, while Section 4.C characterizes the dependency of expected utilities on the information structure. Section 5 discusses some alternative specifications and interpretations of the model. Section 6 concludes. Proofs are relegated to the Appendix.

2. Motivating Studies and Related Literature

2.1. Psychology. A large body of literature within psychology illustrates a phenomenon termed *confirmatory bias*, people's tendency to interpret new evidence in ways that confirm their current beliefs. To give a few examples, in an important early paper, Oskamp (1965) illustrated the inclination of clinical psychologists to become increasingly confident in their analysis of a case even when additional information does not increase accuracy. 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) illustrated the tendency of football viewers to read players' actions as much more aggressive and deserving of penalty when they are preconceived as "bad," "mean," and "aggressive," as manipulated by the use of different uniform colors.[2]

Another heavily studied effect is *cognitive dissonance*, 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

[2] In fact, there is some evidence suggesting that this phenomenon is hard-wired. The McGork Effect in linguistics refers to the observation that when subjects are shown a video of a person saying d sounds (lips separated), and the accompanying soundtrack is that of b sounds (lips closed), subjects report hearing g (as in "great") sounds. This effect vanishes immediately when subjects are asked to shut their eyes. Thus, it is as if the brain goes automatically into a process of resolving an inconsistency between visual and aural stimuli (see Fromkin (2000)).

people can be expected to move their beliefs into line with their behavior. As an illustration, students who were paid very little to perform a boring task described it to their classmates as relatively interesting. Similarly, students who were asked to give a random ideological speech were more likely to change their view in accordance with their speech the lower their initial pay was. In a different context, Glass (1964) reported that students who were asked to give electrical shocks to victims subsequently lowered their opinions of their victims (see Aronson (1969) and Nisbett and Ross (1991) for an overview of related experiments).

In a similar vein, psychologists have identified the prevalence of *motivated reasoning*, people's propensity to reason (by ways of effectively attending only to some of the available information) in a way that supports their subjectively favored propositions (see, e.g., Dawson et al. (2004), Kruglanski (1999), and Kunda (1990)).

2.2. Political Science. 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.

This type of study has been pursued over the years and gained continuous verification. For example, 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). Graber (1984) illustrated people's excess readiness to process political information consistent with their initially held beliefs. Severin and Tankard (2000) overview a variety of evidence dealing with people's desires for consistency and how they affect their search and interpretation of political news.

2.3. Organizational Behavior. In the organizational behavior literature the term *escalation of commitment* is used for a phenomenon akin to cognitive dissonance. 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. One of the first studies illustrating this phenomenon was done by Staw (1976). Staw demonstrated experimentally that subjects were far more inclined to continue investing in a failing project when they were responsible for its earlier funding decision than when another financial officer was. More recently, McCarthy et al. (1993) illustrated that entrepreneurs who started their own businesses invested more than those who bought businesses from others. Schoorman (1988) showed that employees' performance evaluations by supervisors were affected by whether the supervisors had hired the employees originally or not. Staw et al. (1997) demonstrated that banks' tendency to write off bad loans is correlated with managerial turnover. Camerer and Weber (1999) provide robustness of the phenomenon to several alternative explanations.

2.4. Economics. In economics, Akerlof and Dickens (1982) are probably the first to propose a model of motivated reasoning in the context of workers in hazardous occupations. They suggested that cognitive dissonance may drive workers to underestimate the risk they are exposed to and consequently underinvest in protective equipment.

Several papers have incorporated the effects of beliefs on utility derived from self assessment or "ego". Bodner and Prelec (2001), Koszegi (1999), and Akerlof and Kranton (2000) introduced the ideas of self-signaling and identity. These authors considered agents who value their beliefs about themselves and thus choose actions (but not beliefs) that not only maximize some instrumental utility, but also a utility that reflects their self-perceptions. Compte and Postlewaite (2003) considered a related model in which agents' probability of succeeding in an activity increases with their confidence, their perceived probability of success. Their agents selectively remember more previous successes than previous failures, which leads to the agents' overconfidence.

Other types of models concerning the direct effect of beliefs have also been explored. For example, Caplin and Leahy (2001) pointed formally to the effects of anticipation, arising from beliefs concerning future unraveling of uncertainty, on risk aversion. Eyster (2003) explored a model in which regret, arising from beliefs concerning the suboptimality of past actions, directly affects utility. Karlsson et al. (2004) studied the strategic choice of information sources (and the likelihood of strategic ignorance) by investors who are sensitive to regret.

Rabin and Schrag (1999) presented a reduced form model of confirmatory bias in which, with some exogenously given probability (endogenized in the current paper), people misread signals that contradict their current beliefs. The underlying message of Rabin and Schrag (1999)'s paper was that agents prone to confirmatory bias would exhibit overconfidence.

Benabou and Tirole (2002) and Carrillo and Mariotti (2000) considered agents who are time inconsistent and showed that when contemporary information has externalities on future outcomes (market perceptions or consumption), agents may strategically choose to ignore or distort some of the information available to them in each period.

In strategic situations, Geanakoplos et al. (1989) provided a general model in which players' full hierarchies of beliefs enter their utility functions.

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.

Eliaz and Spiegler (2004) illustrate the limitations of the standard vNM model in capturing a set of well-documented anomalies pertaining to information choices. Their work is suggestive of the potential merits of the line of inquiry pursued in the current paper.

Allowing agents to directly choose their beliefs appears in two recent papers. Bracha (2004) suggests a static model in which beliefs are an equilibrium outcome of a game between two internal accounts, choosing beliefs and actions, where the account choosing beliefs is subject to affective motivations. Her model provides another explanation for over- and under-optimism.

Possibly the closest model to the one presented here is that of Brunnermeier and Parker

(2004). They consider an intertemporal consumption model in which agents trade off the benefits of optimism against the negative effects that biased expectations have on the quality of their decision-making. The agents are free to choose their priors, but are constrained to update in a Bayesian manner throughout the (consumption) process. The current paper is complementary in that it explores the effects of unconstrained belief choices and generates results concerning the information structures that agents would or should choose.

3. Setup

This section spells out the model analyzed in this paper. As with many models, there are numerous alternative specifications one could entertain. Since the goal of the current investigation is to simply explore the potential implications of having beliefs influence preferences directly, I will turn to describe the main results immediately after laying out the model. Nonetheless, the reader is directed to Section 5 for an elaborate discussion of other plausible assumptions and their possible effects on the paper's results.

A. Underlying Framework

I consider a world with two possible states $\omega \in \quad = \{L, R\}$, with a-priori probability $\Pr(\omega = L) = 1 - \Pr(\omega = R) = p$, which is transparent to the decision maker. Without loss of generality, I assume that $p > \frac{1}{2}$. That is, L is more likely to be the state of the world.

A decision-maker faces a repeated, three period, choice problem. Time is indexed by $t = 0, 1, 2$.

At the beginning of each period t, the agent receives a signal $s_t \in \{L, R\}$ of accuracy $q_t \in (\frac{1}{2}, 1)$. That is, $\Pr(s_t = \omega \mid \omega) = q_t$. These signals are conditionally independent across time.

At each stage, after having observed that period's signal, the agent tries to guess the state of the world. That is, the agent's set of actions is $A_t \equiv A \equiv \quad = \{L, R\}$ for all t.

The agent's belief at each period t, denoted by μ_t, is her probability assessment of the state of the world ω being L. The prior p serves as the agent's belief at the outset of the game, i.e. $\mu_0 = p$.

The agent's overall utility is comprised of two types of indices. The instantaneous *instrumental utility*, denoted by $u \equiv u(a, \omega)$ will mirror the standard considerations in intertemporal decision problems. The instantaneous *belief utility*, denoted by $v \equiv v(\mu, \mu')$ will capture the agent's taste for consistency. Using these two building blocks, the agent's *ex-post realized utility* at the end of the process is specified by

$$U = \sum_{\tau=0}^{2} \delta^{\tau-t}[u(a_\tau, \omega) + \gamma v(\mu_{\tau-1}, \mu_\tau)],$$

where δ denotes the agent's discount factor (which will usually be assumed to equal 0 or 1) and the parameter $\gamma \geqslant 0$ denotes the weight the agent puts on her belief utility relative to the instrumental utility. For instance, if $\gamma = 0$, the agent is a standard decision maker who considers only her instrumental utility. In the other extreme, as γ grows infinitely large, the agent cares only about the consistency aspect of her decision making.[3]

The agent's *objective function at time t* corresponds to her utility specification and is given by:

$$E_t \left\{ \sum_{\tau=0}^{2} \delta^{\tau-t}[u(a_\tau, \omega) + \gamma v(\mu_{\tau-1}, \mu_\tau)] \right\}.$$

B. Behavior

In order to specify the decision-maker's choices, one needs to specify the beliefs the agent holds over the state of the world, as well as over her future behavior, at each stage.

B.1. Beliefs about the state of the world

At the beginning of period t, given her currently held belief μ_{t-1}, the agent can calculate the Bayesian posterior, based on μ_{t-1} and the observed signal s_t, which will be denoted by $\mu_t^B \equiv \mu_t^B(\mu_{t-1}, s_t)$. Specifically, if the accuracy of the signal s_t is q_t then:

[3] An axiomatic foundation for generalized discounted utility functional forms which have both actions and beliefs as arguments is provided in Yariv (2001).

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

B.2. Beliefs about future selves

As described, the agent at time t puts a relative weight of γ on her belief utility. When contemplating her choice, the agent perceives her future self to put a weight of $\tilde{\gamma}$, $0 \leqslant \tilde{\gamma} \leqslant \gamma$, at any date $\tilde{t} > t$. Thus, I allow an agent to entertain the thought that in the future she will put more weight on the instrumental (standard) aspect of her decision problem than in the present.

B.3. The Choice Problem

At each stage, beliefs are formed before actions are chosen. Thus, the agent chooses an action consistent with her adopted beliefs in the sense that the action maximizes her expected utility conditional on the belief she would report if pressed to do so.[4]

The agent chooses a belief to balance the trade-off between her instrumental and belief utility. Formally, the agent whose decision process will be analyzed throughout the paper is defined as follows.

Definition (**$(\gamma, \tilde{\gamma})$- *consistency*):** *A $(\gamma, \tilde{\gamma})$- consistent agent is one who chooses, at each period $t = 1, 2$, a belief μ_t so that:*

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

where:

1. $a(\mu) \in \arg\max_{a'} E_\mu u(a', \omega)$;
2. *for all* $\tilde{t} > t, \tilde{t} \leqslant 2, s_{\tilde{t}},\ \mu_{\tilde{t}} \in \arg\max_{\mu_{\tilde{t}}} E_{\mu_{\tilde{t}}^B} \sum_{\tau=\tilde{t}}^{T} \delta^{\tau-t} [u(a(\mu_\tau), \omega) + \tilde{\gamma} v(\mu_{\tau-1}, \mu_\tau)]$.

and an action $a_t(\mu_t) \in \arg\max_{a'} E_{\mu_t} u(a', \omega)$.

[4]This assumption, while admittedly strong, is in line with the extensive literature on cognitive dissonance. It seems to be particularly realistic when agents have to justify their actions (see Tetlock et al. (1989)).

A $(\gamma, \tilde{\gamma})$- consistent agent essentially uses backwards induction in period 1 in order to deduce what belief her period 2 self will choose (when the expected weight on the belief utility is $\tilde{\gamma}$) and optimizes her choice of period 1 belief given that prediction.[5]

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

Examples

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:

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

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.

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

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

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

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

[5] Note that if the definition were to be extended to a longer horizon game, one would need to specify beliefs over the beliefs of future selves (concerning the weight put on the belief utility), beliefs over the beliefs over the beliefs of future selves, etc. Assuming that the weight $\tilde{\gamma}$ is common knowledge amongst all future selves simplifies the analysis of longer horizon models significantly.

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:

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

C. Utility Specification

The paper explores the effects of beliefs entering the utility function directly by concentrating on a particular example of instrumental and belief utilities (see discussion in Section 5).

Instrumental utility takes the form of

$$u(a, \omega) = \begin{cases} 1 & a = \omega \\ 0 & a \neq \omega \end{cases} .$$

In words, the agent gets one instrumental util when her guess of the state is correct and zero utils otherwise.

By definition, a $(\gamma, \tilde{\gamma})-$ consistent agent chooses an action consistent with her chosen instantaneous belief. Hence, the agent's decisions satisfy:

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

Inspired by the extensive psychology literature (see Griffin and Tversky (1992) and references therein), 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. Furthermore, the agent's well-being increases with the strength of this support. In other words, the agent welcomes more confidence that confirms her priors. Formally, the belief term at stage $t+1$ is given by:

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

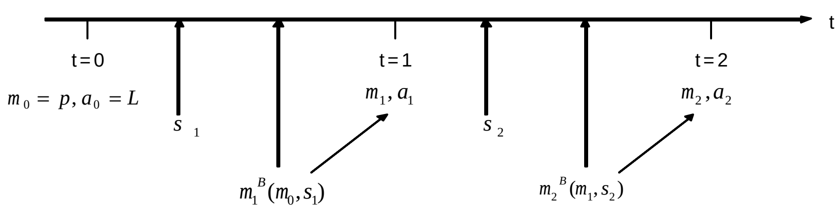

Figure 1: Timeline

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

Figure 1 summarizes the timeline of the environment.

Note that for the above specification, whenever $\gamma \geqslant 2$, the agent is best off choosing $\mu_1 = \mu_2 = 1$ and the instrumental utility plays no role in the analysis. In what follows, I will therefore assume that $\gamma < 2$. Furthermore, since the interest of the current analysis lays in the case of $\gamma > 0$, I will in fact suppose that $\gamma \in (0, 2)$.

Stage 0 actions are effectively exogenous once the given prior μ_0 is set at p. They are specified in the model only for the sake of presentation simplicity.

4. Main Results

The main results are of two types. The first refers to the kind of beliefs agents adopt. I give conditions for agents being overconfident and underconfident. These results are summarized in Subsection A. The second type of results pertains to the levels of expected utility an agent will experience, as is perceived by the agent, and by an omniscient entity who understand the agent's present and future taste for consistency, at the outset of the process, as a function of the amount and quality of the information provided. These results are summarized in Subsection B. Since actions are strongly tied to the beliefs agents hold, a taste for consistency may have observable implications on the action choices. Characterization of conditions that lead to excess persistence and excess volatility of choices (relative to those corresponding to an instrumental utility maximizer) are presented in Subsection C. The proofs of the Propositions stated in this section are relegated to the Appendix.

A. Choice of Beliefs

This section describes the equilibrium choices of beliefs and illustrates the cases under which a $(\gamma, \tilde{\gamma})$ - consistent agent exhibits over and under-confidence relative to the beliefs a pure Bayesian updater would hold.

The first proposition identifies a subset of beliefs the agent may hold in equilibrium, depending on the parameters of the problem:

Proposition 1 (equilibrium beliefs): *In equilibrium, first period beliefs of a $(\gamma, \tilde{\gamma})$ - consistent agents take one of four possible values: $\mu_1^e \in \{\tilde{\mu}_1, \frac{1}{2}, \bar{\mu}_1, 1\}$, where $\tilde{\mu}_1 = \min\{\frac{(2+\tilde{\gamma})(1-q_2)}{2-\tilde{\gamma}(2q_2-1)}, \frac{1}{2}\}$ and $\bar{\mu}_1 = \max\{\frac{(2-\tilde{\gamma})q_2}{2-\tilde{\gamma}(2q_2-1)}, \frac{1}{2}\}$. Second period equilibrium beliefs take two possible values: $\mu_2^e \in \{0, 1\}$.*

Intuitively, in period 1 the agent can potentially choose different beliefs upon observing $s_1 = L$ and $s_1 = R$. Consider the case of $s_1 = L$. The agent certainly prefers to choose beliefs that justify the action $a_1 = L$. In doing so, the agent has to decide whether to choose beliefs that do allow her future self to switch actions or not. If she chooses beliefs that do not allow her future self to switch actions, she is best off gaining the most possible through her current belief utility by ways of choosing $\mu_1 = 1$. If she chooses beliefs that allow her future self to switch actions, she is best off choosing the highest possible probability assessments of L being the realized state that she predicts would lead to a future belief choice no larger than $\frac{1}{2}$ (justifying $a_2 = R$) when observing the signal $s_2 = R$. The value of this belief ends up being

$\bar{\mu}_1$ (see Appendix for exact calculations), when the future weight put on the belief utility is $\tilde{\gamma}$ (as perceived in period 1). Alternatively, she may choose $\mu_1 = \frac{1}{2}$ that, while causing a greater temporary loss of belief utility, will assure the highest possible gain of belief utility in period 2 (of $\frac{\gamma}{2}$).

In period 2, the agent makes choices based on the parameter γ (rather than the predicted $\tilde{\gamma}$). Since $\frac{(2-\tilde{\gamma})q_2}{2-\tilde{\gamma}(2q_2-1)} > \frac{(2-\gamma)q_2}{2-\gamma(2q_2-1)}$ whenever $\gamma > \tilde{\gamma}$, the agent with belief $\mu_1 = \frac{(2-\tilde{\gamma})q_2}{2-\tilde{\gamma}(2q_2-1)}$ will disregard her signal in period 2 and continue choosing the action L, thereby making $\mu_2 = 1$ the optimal choice. Of course, if $\mu_1 = \frac{1}{2}$, the agent will be best off following her signal in period 2 and choosing the most extreme beliefs: 0 if $s_2 = R$ and 1 if $s_2 = L$. A similar intuition holds for the case of $s_1 = R$.

Note that for $(\gamma, \tilde{\gamma})$- consistent agents the chosen beliefs coincide with the Bayesian assessments conditional on observed signals for a non-generic class of parameters. Furthermore, accuracies of future signals may have an effect on current period's beliefs. That is, beliefs at period 1 may depend on the accuracy of the future signal q_2. This is in contrast to standard learning models (see Fudenberg and Levine (1998)). Intuitively, 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 likely to guess the right state of the world then. Therefore, she may be willing distort her current beliefs to boost her instantaneous (period 1) belief utility.

Figure 2 provides an illustration of the equilibrium belief choices.

The range $\left(\frac{(2+\tilde{\gamma})(1-q_2)}{2-\tilde{\gamma}(2q_2-1)}, \frac{(2-\tilde{\gamma})q_2}{2-\tilde{\gamma}(2q_2-1)}\right)$ corresponds to stage 1 beliefs that assure the agent (with perceived belief utility weight $\tilde{\gamma}$) that she will react to the information received at stage 2. I will henceforth assume that stage 2 signals are sufficiently accurate so that this range is non-empty.[6] This assumption translates to the following condition.

[6] When this range is empty, existence of an optimal strategy becomes problematic for sufficiently high q_1 (such that $\frac{1+\delta}{\delta}(1-\mu_1^B(s_1 = R)) > q_2$). Namely, an agent observing the signal $s_1 = R$ would prefer to choose μ_1 as close to, but strictly lower than, $\frac{1}{2}$.

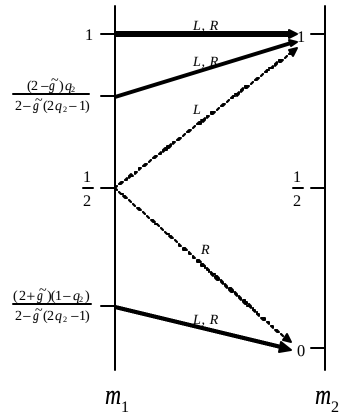

Figure 2: Equilibrium Beliefs

Perceived informative stage 2 signals: $q_2 > \frac{2+\tilde{\gamma}}{4}$.

I classify chosen beliefs according to how they stand relative to the belief an agent with $\gamma = \tilde{\gamma} = 0$ would adopt, i.e. the Bayesian posterior. 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 $\mu > \frac{1}{2}$; or 2. $\mu < \mu^B$ and $\mu < \frac{1}{2}$.

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

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

I will henceforth refer to underconfidence and overconfidence only relative to the beliefs a Bayesian observer would construct.[7] That is, the agent is termed *over(under)-confident in period 1* if the chosen μ_1 is over(under)-confident relative to μ_1^B. The agent is termed *over(under)-confident in period 2* if the chosen μ_2 is over(under)-confident relative to $\mu_2^B(\mu_1^B)$.

In period 2, Proposition 1 implies that the agent always chooses extreme beliefs, $\mu_2^e \in \{0,1\}$ and so is always overconfident.

As for the first period, when the agent cares about consistency, she will generically distort her first period beliefs. From Proposition 1, it follows that first period beliefs are in the set $\{\tilde{\mu}_1, \frac{1}{2}, \bar{\mu}_1, 1\}$. Assume that the first signal is informative. There are three intuitive points to notice. First, if the agent heavily discounts her future payoffs, she is best off maximizing her first period belief utility subject to choosing the optimal instrumental action. In particular, the agent will choose $\mu_1 = 1$ when $s_1 = L$, making her overconfident, and $\mu_1 = \frac{1}{2}$ when $s_1 = R$, making her underconfident. Second, regardless of the discount factor, the agent will choose $\mu_1 = 1$ when $s_1 = L$ if she is willing to ignore the second period's signal, which is the case when $q_2 < \mu_1^B(s_1 = L)$. Similarly, the agent will choose $\mu_1 = \frac{1}{2}$ when $s_1 = R$ if she is keen on using the second period's signal which would intuitively entail $q_2 > 1 - \mu_1^B(s_1 = R)$. Third, regardless of all other parameters, whenever the agent puts sufficiently high weight on her belief utility, she would be best off choosing $\mu_1 = \frac{1}{2}$ when $s_1 = R$. This ensures her the minimal belief utility loss subject to first period instrumental optimality. Furthermore, it assures her a (maximal) gain of $\frac{1}{2}$ on belief utility in period 2. Formally,

Proposition 2 (exaggerated confidence levels): *Assume that* $q_1 > p$.

1. *Whenever* $1 - \mu_1^B(s_1 = R) < q_2 < \mu_1^B(s_1 = L)$, *the* $(\gamma, \tilde{\gamma})$*-consistent agent is overconfident when* $s_1 = L$ *and underconfident when* $s_1 = R$ *in period 1.*
2. *There exists a* $\delta^* \in (0,1)$ *such that as long as* $\delta < \delta^*$, *the* $(\gamma, \tilde{\gamma})$*-consistent agent is overconfident when* $s_1 = L$ *and underconfident when* $s_1 = R$ *in period 1.*

[7] a similar definition appears in the economics literature in, e.g., Rabin and Schrag (1999) and in the psychology literature in, e.g., Griffin and Tversky (1992)

3. *There exists a* $\gamma^* \in (0,2)$ *such that as long as* $\gamma \geqslant \gamma^*$ *the* $(\gamma, \tilde{\gamma})$*-consistent agent is underconfident when* $s_1 = R$ *in period 1.*

The proof, appearing in the Appendix, contains the full characterization of parameters leading to over- and underconfidence following the two possible signals.

B. 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. An agent with an accurate signal can garble her information and imitate an agent with a less accurate signal.

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 consistently), one can determine which structures of information would lead the agent to higher expected utility. I use the term *experienced 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.

Another approach considers the agent's perspective. In a market situation, the agent may 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 beginning of her decision process. 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 $\widetilde{\gamma}$).[8]

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

[8] In this section, I measure both experienced and expected utility from the perspective of the period 1 agent, the period in which effective decisions start taking place. One could replicate the analysis for period 0 (assuming all future selves are believed to have preferences, as well as beliefs about future preferences, characterized by a weight $\tilde{\gamma}$). The analysis relating to experienced expected utility would remain the same. The analysis pertaining to the perceived expected utility is qualitatively similar.

These results provide a potential avenue for testing the theory presented in this paper. The idea that people may prefer less information provides grounds for falsifying the theoretical model presented in this paper. Hysteresis in beliefs and actions (as will be seen below) puts restrictions on the observable choices of information structures.[9]

In what follows, I start by looking at the experienced expected utility. Intuitively, assume that $\widetilde{\gamma}$ 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 Bayesian updating will take place in the second period. However, since the agent actually does care about consistency of beliefs in the second period, she might end up not updating her beliefs when first period beliefs are extreme. 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 decrease 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 present belief - future signals are expected to be so accurate that they will lead to the right choice no matter what current belief is chosen. 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,

[9] This idea is present, to some extent, in the medical literature, which draws a correlation between being in a high risk group for a disease (such as a genetic inclination for certain types of cancer) and reluctance to perform diagnostic tests (see, e.g. Lerman et al. (1999)). Similar examples are given for smokers (see Brock and Balloun (1967)). In addition, the literature in organizational behavior provides some experimental evidence that feedback on decisions may lead managers to avoid information in order not to experience regret (e.g., Larrick and Boles (1995)).

from a Bayesian perspective, when the signal in the second period is lower. This intuition is captured in the following proposition.

Proposition 3 (experienced utility non-monotonicity):

1. *(w.r.t. q_1) For any q_2, there exist $p^* > \frac{1}{2}, \gamma^* > 0$, $\delta^* > 0$ such that for all $p < p^*$, $\widetilde{\gamma} \leqslant \gamma < \gamma^*$, for all $\delta > \delta^*$, expected experienced utility is non-monotonic in q_1.*
2. *(w.r.t. q_2) For any γ, There exist $p^* > \frac{1}{2}, \delta^* > 0$, $\widetilde{\gamma}^* > 0$, and a mapping $q_1^*(p, \delta, \gamma)$* with image in $(\frac{1}{2}, 1)$, *such that for all $p < p^*, \delta < \delta^*, \widetilde{\gamma} \leqslant \min\{\widetilde{\gamma}^*, \gamma\}$,* and $q_1 > q_1^*(p, \delta, \gamma)$, *expected experienced 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, the intrinsic agency problem may lead her to prefer less accurate future signals.

To sharpen intuition, assume, as before, that $\widetilde{\gamma}$ 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, in period 1, that she is more likely to change her beliefs in the second period, even though 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. *(w.r.t. q_1) For all parameters, expected perceived utility is non-decreasing in q_1.*
2. *(w.r.t. q_2) Assume $q_1 > p$. There exists a $\gamma^* > 0$ such that for any δ, there exists $\widetilde{\gamma}^* \leqslant \gamma$ that assures that for all $\widetilde{\gamma} < \widetilde{\gamma}^*$ expected perceived utility is non-monotonic in q_2.*

As it turns out, the results for forward looking agents, agents for whom $\widetilde{\gamma} = \gamma$ 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.[10]

One seemingly peculiar feature of the model is that uninformative signals may increase agents' perceived expected utility. Indeed, even if the decision process ended in period 1, as long as the accuracy of the signal s_1 is lower than that of the prior, p, our agents 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 agents would be willing to purchase such an uninformative signal. The willingness to pay for an uninformative signal is an artifact of the assumption that the agent can change her beliefs only after receiving a signal. One might then wonder whether this is a plausible assumption (as contrasted with, e.g., allowing the agent to spontaneously update her beliefs). In addition, it is natural to question why any event cannot serve as a signal for such an agent. As it turns out,

[10] A preliminary version of the current paper, Yariv (2002), contains a full analysis of the forward looking case.

experimental research suggests that people cannot use absolutely any event to justify their hypotheses.[11]

The model presented here can be thought of as dealing only with signals that have a (latent) proxy for relevance. In other words, signals can be thought of as multi-dimensional, entailing a (binary) dimension that indicates the relevance of each signal to the issue at hand. Signals can be used in a decision process only if they are relevant. The analysis presented in this paper is germane for signals that are relevant for the decision problem at hand.

C. Choice of Actions - Excess Persistence and Volatility

Since actions are strongly tied with beliefs, the persistence of beliefs that consistency implies will be connected with the persistence of actions.

Of course, a standard instrumental utility maximizer may optimally choose identical actions over consecutive periods. For instance, whenever $q_1 < p$, such an agent would optimally choose the same action (L) in the first two periods, periods 0 and 1, with probability 1. If $q_1 > p$, the instrumental utility maximizer would still strictly prefer to follow the period 1 signal and choose L with probability q_1. Similarly, the instrumental utility maximizer would ignore her period 2 signal and choose an action coinciding with that chosen in period 1 whenever $q_2 < \max\{\mu_1, 1-\mu_1\}$, where μ_1 is the prior in the beginning of period 2.

The point of the current section is to illustrate that explicit considerations of beliefs may create excessive persistence and excessive volatility in action choices relative to choices made by standard instrumental utility maximizers. To that effect, we define excess persistence and volatility as follows:

Definition (excess persistence and volatility): *An agent is said to exhibit* excessive persistence (volatility) *at stage t if the probability that $a_t = a_{t-1}$ is greater (smaller) than the corresponding probability prescribed by an instrumental utility maximizer.*

[11] Kunda (1990) notes that 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."

Note that persistence and volatility do not fully capture the instrumental optimality of the actions chosen. These notions relate only to *changes* in choices.

Intuitively, assume the agent places a positive weight on her second period behavior and consider a situation in which the first period's signal indicates that the state is L. As long as $\delta < 1$, the agent would prefer to postpone any loss in belief utility and will prefer to choose $\mu_1 > \frac{1}{2}$ when $s_1 = L$. Since all such choices lead the agent to a choice of L in period 2 (regardless of s_2), the agent will exhibit persistence across periods 1 and 2. Whenever parameters are such that an instrumental maximizer would conceivably change actions after $s_1 = L$ (namely, when $q_2 > \mu_1^B(s_1 = L)$), the agent will be excessively persistent in period 2.

Excess persistence in both periods requires that an agent observing a signal $s_1 = R$ would still prefer to choose $\mu_1 > \frac{1}{2}$. This entails a high regard for consistency and a sufficiently low discount factor. Indeed, as δ becomes smaller, the agent cares less about her future choices and so is willing to gain on her immediate belief utility at the expense of making wrong future choices with higher probabilities. As it turns out, for sufficiently low δ, there exists a positive measure of parameters (in particular, sufficiently high $\gamma's$ and any corresponding $\tilde{\gamma} \leqslant \gamma$) for which an instrumental maximizer would switch actions with positive probability, but a $(\gamma, \tilde{\gamma})$-consistent agent would always choose L. These cases correspond to the agent being excessively persistent in both periods.

Since the agent is inclined to choose beliefs closer to her prior $p > \frac{1}{2}$, excess volatility in period 1 never occurs in this setup. Excess volatility in period 2 may occur only when period 1 beliefs are chosen to allow the agent a later switch of actions, namely when $\mu_1 = \frac{1}{2}$. If $\mu_1 = \frac{1}{2}$ is chosen when $s_1 = L$ by a $(\gamma, \tilde{\gamma})$-consistent agent, the agent is willing to forgo belief utility (which could have been achieved by choosing $\mu_1 = 1$) in order to gain flexibility in period 2. In that case, an instrumental utility maximizer, who does not experience this belief cost, would surely respond to the second period signal as well. Consequently, excess volatility may occur only when the $(\gamma, \tilde{\gamma})$-consistent agent chooses $\mu_1 = \frac{1}{2}$ (and $a_1 = R$) after observing $s_1 = R$, while the instrumental utility maximizer would choose R in periods 1 and 2. Note that choosing $a_1 = R$ and assuring that $a_2 = R$ entails a choice of $\mu_1 = \tilde{\mu}_1$, leading to overall belief utility of $\gamma\left[(\tilde{\mu}_1 - p) + \delta\tilde{\mu}_1\right]$. On the other hand, choosing $\mu_1 = \frac{1}{2}$ yields overall belief utility of $\gamma[(\frac{1}{2} - p) + \delta\frac{1}{2}]$. Hence, the gain in belief utility of choosing $\mu_1 = \frac{1}{2}$ over choosing $\mu_1 = \tilde{\mu}_1$ is $(1+\delta)(\frac{1}{2} - \tilde{\mu}_1)$, which is increasing in δ. For sufficiently large δ, there are parameters for which this gain outweighs the potential instrumental loss leading to excess volatility.

Proposition 5 formalizes the above intuition. In Proposition 5 we concentrate on the case in which the first signal is informative, $q_1 > p$, so that an instrumental utility maximizer would follow the signal in period 1 and choose each action L and R with a positive probability.[12]

Proposition 5 (occurence of persistence and volatility): *Assume $q_1 > p$.*

1. *For all $\delta < 1$, the $(\gamma, \tilde{\gamma})$-consistent agent exhibits excess persistence in period 1 for a positive measure of all other parameters. Furthermore, there exists $\delta^* > 0$ such that for all $\delta < \delta^*$ the $(\gamma, \tilde{\gamma})$-consistent agent exhibits excess persistence in both periods for a positive measure of all other parameters.*

2. *For all parameters, a $(\gamma, \tilde{\gamma})$-consistent agent is never excessively volatile in period 1. There exists $\delta^{**} \in [0, 1)$ such that for all $\delta > \delta^{**}$, there exists a positive measure of all other parameters for which the $(\gamma, \tilde{\gamma})$-consistent agent is excessively volatile in period 2.*

The proof of Proposition 5 contains the full necessary and sufficient conditions for excess persistence in either period and excess volatility in period 2.

Excess persistence and volatility are phenomena observed in the field (for an overview, see, e.g., Constantinides (1990)). Proposition 5 explains these observations as arising from direct care for consistency. Furthermore, it suggests a correlation between agents' propensity for persistence and their intertemporal discount factors.

[12] The analysis of $q_1 < p$ is similar in nature, though excess persistence and volatility are plausible only in period 2.

5. DISCUSSION OF THE MODEL

A. On the Belief Utility

The specification of belief utility allows me to gain insights on the implications of having a taste for consistency on agents' decision-making process. The use of linear directional confidence serves as a benchmark for this kind of analysis. Of course, there are many alternative ways to model belief utility. For example:

1. A prominent consistency theory in the psychology literature is that of disappointment and perceived utility maintenance (see, e.g., 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:

$$\begin{aligned} v(\mu_t, \mu_{t+1}) &= dissapointment(a_t(\mu_0), \mu_{t+1}) = \\ &\quad w((\mu_{t+1} - \mu_t) I_L(a_t(\mu_t)) + (\mu_t - \mu_{t+1}) I_R(a_t(\mu_t))) \end{aligned}$$

   where $I_j(k) = \{ \begin{array}{ll} 1 & j = k \\ 0 & j \neq k \end{array}$.

   Taking a linear representation: $w(x) = x$, we get:

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

   For $\mu_t = \frac{1}{2}$ we can arbitrarily choose $v(\mu_t, \mu_{t+1}) = \mid \mu_{t+1} - \mu_t \mid$ and v is, in fact, identical to the linear directional confidence belief utility.

2. An alternative form of belief utility is *regret utility* (see Loomes and Sugden (1982)). 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:

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

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

when the regret function R is taken to be linear: $R(x) = x$. Here previous beliefs play a role in the belief utility only to the extent that they determine the previous period's action. Regret utilities lead to similar qualitative results and are not analyzed here.

B. On the Updating Process

The model presented in this paper assumes that agents use $\mu_t^B \equiv \mu_t^B(\mu_{t-1}, s_t)$ when choosing their beliefs and actions. That is, agents calculate the posterior accurately and then choose a belief and a corresponding action that balance the trade-off between the instrumental utility and the belief utility, solving a problem of the sort $\max_{\mu_t} E_{\mu_t^B} U_t$. This assumption is made to minimize the distance between the current model and the standard one. Once preferences are specified, the agents maximize their well-being using all the information and statistical tools they have at hand. Nonetheless, I view this as a strong assumption that serves as a benchmark, the modifications of which are left for future research.

C. On Finite Memory

Implicit in the formulation is the assumption that at each stage t, agents remember only their previous action a_{t-1} and belief μ_{t-1}.[13] In the standard framework in which beliefs always coincide with the Bayesian posteriors, i.e., $\mu_t = \mu_t^B$, previous period's beliefs are sufficient statistics for the information accumulated. In the current setup, however, remembering the entire set of signals could make a difference to the agent's decisions.

The assumption that agents do not remember the full history of signals and beliefs has a few justifications in this framework:

[13] In fact, since actions are required to be consistent with beliefs, the agent can deduce her past actions even if she only remembers her past beliefs.

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

2. The assumption that agents remember previous beliefs rather than previous signals and corresponding accuracies is grounded in 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)).

3. Conceptually, allowing the agents to remember their full history of choices would introduce concerns regarding the possible manipulation of the agents' memories. Investigating how agents edit their memories (encompassing activities such as selective memory and hind-sight biases) seems necessary before introducing larger blocks of memory into the current model. I find this an important avenue for extensions of the model.

D. 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 (see Bodner and Prelec (2001), Koszegi (1999), and Akerlof and Kranton (2000)).

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.

The current paper does not contribute to this debate. In particular, I do not attempt to connect the results presented here with ideas about signaling of competence (to oneself or to others).

6. Concluding Comments

A large body of experimental and empirical work indicates that people exhibit cognitive biases implying a taste for consistency. In broad terms, this paper's contribution is in providing an exploratory 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, as well as excess stickiness in action choices, and predicts agents' preference for less accurate signals in some circumstances. The results are consistent with evidence concerning observed overconfidence of financial investors, selective exposure to political information, escalation of commitment of business managers, and more.

Given the abundant literature on motivated reasoning extant in the psychology literature, it is important to scrutinize the contribution of a modeling endeavor such as the one presented here. Specifically, in concluding the paper it is worthwhile noting a few virtues of the current paper over the existing psychology literature: 1. the model allows to link several phenomena, such as confirmatory bias and cognitive dissonance, under the same umbrella of preferences that depend on beliefs per-se; 2. the model allows to organize predictions across choice problems - for example, the fact that an agent is over-confident in one situation has implications on how she will behave in others; 3.writing down the model helps clarify what are the important parameters for predicting behavior (such as γ and $\widetilde{\gamma}$) as well as allude to phenomena that have not been tested for before in economic settings (such as the desire for less accurate or less information); but 4. just like any model, this framework may not always be appropriate. Future studies would be helpful in determining when this setup can be more useful than the standard one.

Even if one accepts the underlying framework suggested in this paper, there is still a lot of ground for more work. In what follows I suggest a few directions for future inquiry.

To begin, the model presented here is a benchmark model in a few respects. For example, it would be useful to explore longer decision process and the implication of different information statistics than solely previous period's belief (thereby entailing a more complex model of memory). It may also be useful to relax the assumption that agents maximize their objective function with respect to the Bayesian posterior, by 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, etc.

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.

The current paper concentrated on a single agent decision game. It may prove useful to embed a model of the sort presented here in a strategic setting. Such an analysis may provide a foundation for self-confirming equilibria (Fudenberg and Levine (1993)). Indeed, if all 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. It seems natural to conjecture that the speed of convergence would be related to the difference between the actual parameter γ and the perceived $\widetilde{\gamma}$, but analytical results are yet to be derived.

Appendix - Proofs of Main Results

The following Lemma is of use in all subsequent proofs. In particular, the proof of Proposition 1 follows directly from it.

Lemma: *The choice of* μ_1 *and* a_1 *at stage 1 is chosen to maximize* $E_{t=1,\mu_1^B|\mu_1}U = p + \delta V(a_1,\mu_1)$*, where*

$$V(a_1,\mu_1) = \begin{cases} (1+\delta)\mu_1^B + \gamma[\mu_1(1-\delta)+\delta-p] & \mu_1 > \bar{\mu}_1 \\ \left[\mu_1^B + \gamma(\mu_1 - p)\right] + \delta\left[q_2 - \gamma\left(\frac{q_2+\mu_1^B}{2} - 1 + \mu_1 - q_2\mu_1^B\right)\right] & \frac{1}{2} < \mu_1 \leqslant \bar{\mu}_1 \\ \left[\mu_1^B + \gamma(\frac{1}{2} - p)\right] + \delta\left(q_2 + \frac{\gamma}{2}\right) & \mu_1 = \frac{1}{2}, a_1 = L \\ \left[1 - \mu_1^B + \gamma(\frac{1}{2} - p)\right] + \delta\left(q_2 + \frac{\gamma}{2}\right) & \mu_1 = \frac{1}{2}, a_1 = R \\ \left[1 - \mu_1^B + \gamma(\mu_1 - p)\right] + \delta\left[q_2 + \gamma\left(\frac{q_2+\mu_1^B-1}{2} + \mu_1 - q_2\mu_1^B\right)\right] & \tilde{\mu}_1 < \mu_1 < \frac{1}{2} \\ (1+\delta)(1-\mu_1^B + \gamma\mu_1) - \gamma p & \mu_1 \leqslant \tilde{\mu}_1 \end{cases},$$

where $\bar{\mu}_1 = \frac{(2-\tilde{\gamma})q_2}{2-\tilde{\gamma}(2q_2-1)}$ *and* $\tilde{\mu}_1 = \frac{(2+\tilde{\gamma})(1-q_2)}{2-\tilde{\gamma}(2q_2-1)}$.

Proof: Using backward induction.

At stage 2, assuming first that $\mu_1 > \frac{1}{2}$.

If $q_2 \leqslant \mu_1$ then $\mu_2^B \geqslant \frac{1}{2}$ for all s_2 and $\mu_2 = 1$ leading to a choice of $a_0 = a_1 = a_2 = L$.

If $q_2 > \mu_1$ then $\mu_2^B(s_2 = R) < \frac{1}{2}$ so the agent is expected to make a 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,

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

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

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

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

If $q_2 > 1-\mu_1$ then $\mu_2^B(s_2 = L) > \frac{1}{2}$ and the agent is expected to make a 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 $\frac{(1-\mu_1)(1-q_2)}{\mu_1 q_2 + (1-\mu_1)(1-q_2)} + \tilde{\gamma}\mu_1$.

Therefore,

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

(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 $\mu_1 \geqslant \frac{1}{2}$, there are two cases to consider:

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

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

Assume now that $\mu_1 < \frac{1}{2}$ then 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 stage 1 expected utility is:

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

The assumption that $q_2 > \frac{2+\tilde{\gamma}}{4}$ guarantees that $\frac{(2+\tilde{\gamma})(1-q_2)}{2-\tilde{\gamma}(2q_2-1)} < \frac{1}{2}$ and $\frac{(2-\tilde{\gamma})q_2}{2-\tilde{\gamma}(2q_2-1)} > \frac{1}{2}$.The above equalities, together with the constrained consistency between a_1 and μ_1, generate the Lemma's claim (where endpoints of belief ranges are potentially chosen only if they maximize the beliefs in that range as payoffs increase with the level of beliefs in each range). ■

Proof of Proposition 1: From the Lemma, since $V(\mu_1, a_1)$ is monotonic in μ_1 within each of the specified intervals, the agent will always choose $\mu_1 \in \left\{\frac{(2+\tilde{\gamma})(1-q_2)}{2-\tilde{\gamma}(2q_2-1)}, \frac{1}{2}, \frac{(2-\tilde{\gamma})q_2}{2-\tilde{\gamma}(2q_2-1)}, 1\right\}$. Furthermore, since $\frac{(2+\tilde{\gamma})(1-q_2)}{2-\tilde{\gamma}(2q_2-1)} < \frac{(2+\gamma)(1-q_2)}{2-\gamma(2q_2-1)}$ and $\frac{(2-\tilde{\gamma})q_2}{2-\tilde{\gamma}(2q_2-1)} > \frac{(2-\gamma)q_2}{2-\gamma(2q_2-1)}$, in equilibrium, $\mu_2 \in \{0, 1\}$. ■

Proof of Proposition 2: In period 1, assuming the agent chooses $\bar{\mu}_1$ when indifferent between $\bar{\mu}_1$ and 1, $\frac{1}{2}$ when indifferent between $\bar{\mu}_1$ and $\frac{1}{2}, \tilde{\mu}_1$ and $\frac{1}{2}$, or 1 and $\frac{1}{2}$ (thereby choosing less extreme beliefs), the agent is overconfident when $\mu_1^B \geqslant \frac{1}{2}$ if either:

- $\min\left\{\gamma\left[(1-\bar{\mu}_1) + \delta\left(\frac{q_2+\mu_1^B}{2} - 1 + \bar{\mu}_1 - q_2\mu_1^B\right)\right], \frac{(1-\delta)\gamma}{2}\right\} > \delta\left(q_2 - \mu_1^B\right)$ in which case $\mu_1 = 1 > \mu_1^B$; or
- $$\begin{aligned} &\gamma\left[(1-\bar{\mu}_1) + \delta\left(\frac{q_2+\mu_1^B}{2} - 1 + \bar{\mu}_1 - q_2\mu_1^B\right)\right] \leqslant \delta\left(q_2 - \mu_1^B\right), \\ &\gamma\left[\left(\bar{\mu}_1 - \tfrac{1}{2}\right) - \delta\left(\frac{q_2+\mu_1^B-1}{2} + \bar{\mu}_1 - q_2\mu_1^B\right)\right] > 0, \\ &\text{and } \bar{\mu}_1 > \mu_1^B, \end{aligned}$$
  in which case $\mu_1 = \bar{\mu}_1 > \mu_1^B$.

The agent is overconfident when $\mu_1^B < \frac{1}{2}$ if

$$\begin{aligned} &(1+\delta)\gamma\left(\tfrac{1}{2} - \tilde{\mu}_1\right) < \delta(1 - q_2 - \mu_1^B) \\ &\text{and } \tilde{\mu}_1 < \mu_1^B, \end{aligned}$$

in which case $\mu_1 = \tilde{\mu}_1 < \mu_1^B$.

Similarly, the agent is underconfident in period 1 when $\mu_1^B \geqslant \frac{1}{2}$ if either:

- $$\begin{aligned} &\gamma\left[(1-\bar{\mu}_1) + \delta\left(\frac{q_2+\mu_1^B}{2} - 1 + \bar{\mu}_1 - q_2\mu_1^B\right)\right] \leqslant \delta\left(q_2 - \mu_1^B\right), \\ &\gamma\left(\bar{\mu}_1 - \tfrac{1}{2}\right) > \delta\gamma\left(\frac{q_2+\mu_1-1}{2} + \bar{\mu}_1 - q_2\mu_1^B\right), \\ &\text{and } \bar{\mu}_1 < \mu_1^B; \text{ or} \end{aligned}$$
- $$\begin{aligned} &\frac{(1-\delta)\gamma}{2} \leqslant \delta\left(q_2 - \mu_1^B\right), \\ &\gamma\left(\bar{\mu}_1 - \tfrac{1}{2}\right) \leqslant \delta\gamma\left(\frac{q_2+\mu_1-1}{2} + \bar{\mu}_1 - q_2\mu_1^B\right), \\ &\text{and } \tfrac{1}{2} < \mu_1^B; \text{ or} \end{aligned}$$

The agent is underconfident when $\mu_1^B < \frac{1}{2}$ if either:

- $(1+\delta)\gamma\left(\frac{1}{2}-\tilde{\mu}_1\right) \geqslant \delta(1-q_2-\mu_1^B)$, so that $\mu_1 = \frac{1}{2} > \mu_1^B$; or
- $(1+\delta)\gamma\left(\frac{1}{2}-\tilde{\mu}_1\right) < \delta(1-q_2-\mu_1^B)$ and $\tilde{\mu}_1 > \mu_1^B$.

In particular, when $q_1 > p$ (the first signal is informative), $\mu_1^B(s_1 = L) > \frac{1}{2}$ and $\mu_1^B(s_1 = R) < \frac{1}{2}$. Now,

1. Whenever $1 - \mu_1^B(s_1 = R) < q_2 < \mu_1^B(s_1 = L)$, the agent chooses $\mu_1 = 1$ when observing $s_1 = L$ and $\mu_1 = \frac{1}{2}$ when observing $s_1 = R$, which produces the result.
2. For sufficiently low δ, the agent's equilibrium choices are the same as in 1. above and the result follows.
3. For sufficiently high γ, as long as $\tilde{\mu}_1 \leqslant \mu_1^B$, the agent optimally chooses $\mu_1 = \frac{1}{2}$ upon observing $s_1 = R$ for all δ, and the agent is underconfident when observing R in period 1. If $\tilde{\mu}_1 > \mu_1^B$, for sufficiently high δ, the agent will choose $\mu_1 = \tilde{\mu}_1$ upon observing $s_1 = R$, which would still make her underconfident. ∎

Proof of Proposition 3.1 (experienced utility non-monotonicity w.r.t. q_1): From the proof of Lemma 1 we see that the value of q_1 affects the choice of $\mu_1 \in \left\{\frac{(2+\tilde{\gamma})(1-q_2)}{2-\tilde{\gamma}(2q_2-1)}, \frac{1}{2}, \frac{(2-\tilde{\gamma})q_2}{2-\tilde{\gamma}(2q_2-1)}, 1\right\}$ (though not the set of values from which μ_1 is chosen). The experienced utility upon a posterior μ_1^B at stage 1 and choice of beliefs $\mu_1 \in \left\{\frac{(2+\tilde{\gamma})(1-q_2)}{2-\tilde{\gamma}(2q_2-1)}, \frac{1}{2}, \frac{(2-\tilde{\gamma})q_2}{2-\tilde{\gamma}(2q_2-1)}, 1\right\}$ is given by $\mathbb{E}_{t=1,\mu_1^B|\mu_1} U^e = p + \delta V^e(a_1, \mu_1)$, where

$$V^e(a_1,\mu_1) = \begin{cases} (1+\delta)\mu_1^B + \gamma[\mu_1(1-\delta)+\delta-p] & \mu_1 > \frac{1}{2} \\ \left[\mu_1^B + \gamma(\frac{1}{2}-p)\right] + \delta\left(q_2 + \frac{\gamma}{2}\right) & \mu_1 = \frac{1}{2}, a_1 = L \\ \left[1-\mu_1^B + \gamma(\frac{1}{2}-p)\right] + \delta\left(q_2 + \frac{\gamma}{2}\right) & \mu_1 = \frac{1}{2}, a_1 = R \\ (1+\delta)(1-\mu_1^B + \gamma\mu_1) - \gamma p & \mu_1 < \frac{1}{2} \end{cases}.$$

Note that an agent restricted to a choice of $\mu_1 \in \left\{\frac{(2+\tilde{\gamma})(1-q_2)}{2-\tilde{\gamma}(2q_2-1)}, \frac{1}{2}, \frac{(2-\tilde{\gamma})q_2}{2-\tilde{\gamma}(2q_2-1)}, 1\right\}$ and maximizing the experienced utility would prefer more accurate first period signals. Furthermore, such an agent would make choices coinciding with those of the $(\gamma, \tilde{\gamma})$-consistent agent in all cases other than those in which the latter chooses $\mu_1 = \frac{(2-\widetilde{\gamma})q_2}{2-\widetilde{\gamma}(2q_2-1)} \in (\frac{1}{2}, 1)$ (in which case $\mu_1 = \frac{1}{2}$ or $\mu_1 = 1$ would be chosen by the experienced utility maximizer). Since we assumed that $\tilde{\gamma} < \gamma_1^* = 4q_2 - 2$ (perceived informative stage 2 signals), $\bar{\mu}_1 > \frac{1}{2}$. Consider now q_1 for which the agent, upon observing the signal L, is indifferent between choosing $\mu_1 = \frac{1}{2}$ and $\mu_1 = \bar{\mu}_1$, but strictly prefers either to a choice of $\mu_1 = 1$, implying a strictly lower experienced utility for a choice of $\mu_1 = \bar{\mu}_1$ than for a choice of $\mu_1 = \frac{1}{2}$. This occurs if q_1 is such that:

$$\mu_1^B = \frac{1 - \delta - 2\bar{\mu}_1 + \delta q_2 + 2\delta\bar{\mu}_1}{\delta (2q_2 - 1)}; \text{ and}$$

$$\frac{2\delta q_2 - \gamma(1-\delta)}{2\delta} > \frac{1 - \delta - 2\bar{\mu}_1 + \delta q_2 + 2\delta\bar{\mu}_1}{\delta (2q_2 - 1)}.$$

As long as $\delta > \frac{4\bar{\mu}_1 - 2}{4\bar{\mu}_1 - 1}$ (which is no smaller than $\frac{2}{3}$ that can be chosen as δ^*), we have $\frac{1-\delta-2\bar{\mu}_1+\delta q_2+2\delta\bar{\mu}_1}{\delta(2q_2-1)} > \frac{1}{2}$.

Furthermore, for $\gamma < \gamma_2^* \equiv 2\frac{(1-\delta)(2\bar{\mu}_1-1)-2\delta q_2(1-q_2)}{(2q_2-1)(1-\delta)}$, and certainly for $\gamma < \gamma^* \equiv \min(\gamma_1^*, \gamma_2^*)$, the above inequality holds. In particular, $\frac{1-\delta-2\bar{\mu}_1+\delta q_2+2\delta\bar{\mu}_1}{\delta(2q_2-1)}$ is in the range $(\frac{1}{2}, 1)$ and whenever $p < p^* \equiv \frac{1-\delta-2\bar{\mu}_1+\delta q_2+2\delta\bar{\mu}_1}{\delta(2q_2-1)}$, there indeed exists $q_1 = q_1^* > \frac{1}{2}$ for which $\mu_1^B(s_1 = L) = \frac{1-\delta-2\bar{\mu}_1+\delta q_2+2\delta\bar{\mu}_1}{\delta(2q_2-1)}$.

For sufficiently small $\varepsilon > 0$, consider the transition from $q_1^* - \varepsilon$ (yielding a choice of $\mu_1 = \frac{1}{2}$ when $s_1 = L$) to $q_1^* + \varepsilon$ (yielding a choice of $\mu_1 = \bar{\mu}_1$ for $s_1 = L$). There are two effects of such an increase in the first period's signal: 1. Conditional on $s_1 = L$, the increase in q_1 creates a discrete reduction in experienced utility bounded above 0; 2. Conditional on $s_1 = R$, the increase in q_1 creates an increase in the experienced utility of ε order of magnitude. The overall effect corresponds to a decrease in experienced utility when ε is small enough. Note that for $q_1 > q_1^*$, experienced utility is monotonic in q_1. The Proposition's claim then follows. ■

Proof of Proposition 3.2 (experienced utility non-monotonicity w.r.t. q_2): Consider the extreme case in which $\widetilde{\gamma} = 0$. As in the proof of the first part of the proposition, I will illustrate the case in which as q_2 increases, the agent switches her first period belief to one which she perceives will allow her to update in period 2, but that, in reality, does not, and consequently yields a lower overall experienced utility. Formally, upon observing $s_1 = L$, the agent is indifferent between choosing μ_1 to be $\bar{\mu}_1 = \frac{(2-\tilde{\gamma})q_2}{2-\tilde{\gamma}(2q_2-1)} = q_2$ or 1, and strictly prefers either to choosing $\mu_1 = \frac{1}{2}$ when

$$q_2 = \frac{2\delta(\mu_1^B - \gamma p - \gamma) + 2\gamma + \delta\gamma\mu_1^B}{2\gamma - 3\delta\gamma + 2\delta(1 + \gamma\mu_1^B)} \equiv \hat{q}_2; \text{ and}$$

$$\mu_1^B > \frac{1}{2}\frac{-\gamma + 2\delta q_2 + \delta\gamma}{\delta}.$$

Algebraic manipulations illustrate that these specifications are meaningful, i.e., $\hat{q}_2 \in (\frac{1}{2}, 1)$ simultaneously with $\mu_1^B(s_1 = L) > \frac{1}{2}\frac{-\gamma+2\delta q_2+\delta\gamma}{\delta}$, when $p < p^* \equiv \frac{2-\delta}{4\delta}$, $\delta < \delta^* \equiv \min\{\frac{2+2\gamma}{1+4q_2+2\gamma}, \frac{2}{3}\}$ (where $\delta < \frac{2}{3}$ assures that $p^* > \frac{1}{2}$), and q_1 sufficiently high so that $\mu_1^B(s_1 = L)$ indeed satisfies the above inequality.

Note that for sufficiently small $\varepsilon > 0$, an increase from $\hat{q}_2 - \frac{\varepsilon}{2}$ to $\hat{q}_2 + \frac{\varepsilon}{2}$ would cause a decrease of

$$[pq_1 + (1-p)(1-q_1)]\gamma(1-\delta)(1 - \hat{q}_2 - \frac{\varepsilon}{2}) - \frac{\varepsilon}{2}M$$

in the expected experienced utility, where M is a finite number. In particular, the experienced expected value decreases between $\hat{q}_2 - \frac{\varepsilon}{2}$ and $\hat{q}_2 + \frac{\varepsilon}{2}$. It is immediate that for either the range $q_2 > \hat{q}_2$ or the range $\frac{1}{2} \leqslant q_2 < \hat{q}_2$, experienced utility is monotonically increasing. The proposition's claim thereby follows from continuity. ■

Proof of Proposition 4.1 (perceived utility monotonicity w.r.t. q_1): Let $r_1 > q_1 > \frac{1}{2}$. 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 :

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

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

From the choice of y,

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

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): Assume $\tilde{\gamma} = 0$. Denote by:

$$\begin{aligned}
\hat{q} &= \frac{2\delta(1-\mu_1^B(s_1=R)) + \gamma(1+\delta)}{2\left(\delta+\gamma+\delta\gamma\right)}; \\
\check{q} &= \frac{\delta(2+\gamma)\mu_1^B(s_1=L) + 2\gamma(1-\delta)}{2\delta\gamma\mu_1^B(s_1=L) + 2(\delta+\gamma) - 3\delta\gamma}; \\
\breve{q} &= \frac{2\delta\mu_1^B(s_1=L) + \gamma(1-\delta)}{2\delta}
\end{aligned}$$

Note that $\hat{q}$ corresponds to the intersection of the the values of $V(a_1, \mu_1)$ when $\mu_1 = \frac{1}{2}$ and $\mu_1 = 1 - q_2$, upon observing $s_1 = R$. Similarly, $\check{q}$ corresponds to the intersection when $\mu_1 = 1$ and $\mu_1 = q_2$, $\breve{q}$ to the intersection when $\mu_1 = 1$ and $\mu_1 = \frac{1}{2}$ $(a_1 = L)$ upon observing $s_1 = L$. Note that whenever $\mu_1^B(s_1 = R) < \frac{1}{2}$ and $\gamma > \frac{4\delta}{2+3\delta}$ (which is no larger than $\frac{4}{5}$ that can be chosen as γ^*), $\min\{\hat{q},\, \check{q},\, \breve{q}\} > \frac{1}{2}$. As long as $q_2 < \min\{\hat{q},$ $\check{q},\, \breve{q}\}$, the agent chooses $\mu_1 = 1$ (and $a_1 = L$) upon observing $s_1 = L$ and $\mu_1 = 1 - q_2$ (and $a_1 = R$) upon observing $s_1 = R$. In particular, perceived utility decreases with q_2 in that range. Note that whenever $q_2 > \hat{q}$ the agent chooses $\mu_1 = \frac{1}{2}$ whenever $s_1 = R$ and Lemma 1 assures that expected perceived utility is monotonically increasing in q_2. The proposition's claim then follows from continuity. ■

Proof of Proposition 5: Since $q_1 > p$, an instrumental utility maximizer would follow her signal in period 1, and

1. Persistence of actions in both periods occurs with positive probability whenever

$\mu_1(s_1 = L) > \frac{1}{2}$, in which case

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

A necessary and sufficient condition is that the agent prefers choosing $\mu_1 = 1$ or $\mu_1 = \bar{\mu}_1 = \frac{(2-\widetilde{\gamma})q_2}{2-\widetilde{\gamma}(2q_2-1)}$ over $\mu_1 = \frac{1}{2}$ when $s_1 = L$ Thus, persistence of actions occurs with positive probability when either $q_2 < \mu_1^B(s_1 = L) + \frac{\gamma(1-\delta)}{2\delta}$ or when $q_2 > \frac{2\mu_1^B(s_1=L)\delta+2\gamma-2\gamma\bar{\mu}_1+3\delta\gamma\bar{\mu}_1-2\delta\gamma}{\delta(2-\gamma+2\gamma\mu_1^B(s_1=L))}$ (note that $\bar{\mu}_1$ depends on q_2). In particular, excessive persistence occurs in period 2 whenever

$$\begin{aligned} q_2 \in & \left(\mu_1^B(s_1 = L), \mu_1^B(s_1 = L) + \frac{\gamma(1-\delta)}{2\delta}\right) \cup \\ & \left(\max\left\{\mu_1^B(s_1 = L), \frac{2\mu_1^B(s_1 = L)\delta + 2\gamma - 2\gamma\bar{\mu}_1 + 3\delta\gamma\bar{\mu}_1 - 2\delta\gamma}{\delta(2 - \gamma + 2\gamma\mu_1^B(s_1 = L))}\right\}, 1\right). \end{aligned}$$

In fact, since the conditions on q_2 that guarantee $\mu_1(s_1 = R) < \frac{1}{2}$ are stronger than those required for $\mu_1(s_1 = L) > \frac{1}{2}$, the latter restriction provides the requirement for excess persistence in period 2. Notice that these conditions imply that as long as $\delta < 1$, there exists a non-empty range of parameters for which the agent is excessively persistent in period 2, as is proposed.

Excess persistence occurs in period 1 as well whenever $\mu_1(s_1 = R) > \frac{1}{2}$, in which case persistence of actions occurs with probability 1:

$$a_1(s_1) = a_2(s_1, s_2) = L \text{ for all } s_1, s_2.$$

For this condition to hold, it suffices to show that the agent chooses $\mu_1(s_1 = R) > \frac{1}{2}$. The agent chooses $\mu_1 = 1$ whenever q_2 satisfies $q_2 < \mu_1^B(s_1 = R) + \frac{4\mu_1^B(s_1=R)+\gamma(1-\delta)-2}{2\delta}$ and $\tilde{\mu}_1 < \frac{1}{1+\delta} - \frac{(1-2\mu_1^B(s_1=R))}{\gamma}$. Similarly, the agent chooses $\mu_1 = \bar{\mu}_1$ whenever the parameters are such that

$$\gamma[(\bar{\mu}_1 - \frac{1}{2}) - \delta\left(\frac{q_2 + \mu_1^B(s_1 = R) - 1}{2} + \bar{\mu}_1 - q_2\mu_1^B(s_1 = R)\right)] > 1 - 2\mu_1^B(s_1 = R)$$

and

$$\begin{aligned} \mu_1^B(s_1 &= R) + \gamma\bar{\mu}_1 + \delta[q_2 - \gamma\left(\frac{q_2 + \mu_1^B(s_1 = R)}{2} - 1 + \bar{\mu}_1 - q_2\mu_1^B(s_1 = R)\right)] > \\ &> (1 + \delta)(1 - \mu_1^B(s_1 = R) + \gamma\tilde{\mu}_1). \end{aligned}$$

In particular, for sufficiently small δ, there is a positive measure of parameters for which the agent chooses $\mu_1 = 1$, while an instrumental utility maximizer would react to her signals in both periods. The proposition's claim then follows.

2. It is straightforward to see that the $(\gamma, \tilde{\gamma})$-consistent agent never chooses $a_1 = R$ when $\mu_1^B > \frac{1}{2}$. Therefore, the agent is never excessively volatile in period 1. The agent chooses a period 2 action which is different than period 1's with positive probability whenever $\mu_1 = \frac{1}{2}$. Whenever $\mu_1^B(s_1 = L) > q_2$, the agent chooses $\mu_1 > \frac{1}{2}$ when $s_1 = L$. Therefore, excess volatility may occur only when the agent chooses $\mu_1 = \frac{1}{2}$ when $s_1 = R$ and $1 - \mu_1^B(s_1 = R) > q_2$. In fact, the choice of $\mu_1 = \frac{1}{2}$ is surely chosen by an agent observing $s_1 = R$ whenever the following three conditions hold:

   - $\frac{\gamma}{2}(1+\delta) + \delta q_2 > \delta(1 - \mu_1^B(s_1 = R)) + (1+\delta)\gamma\tilde{\mu}_1$;
   - $1 + \frac{\gamma}{2}(1 + \delta q_2) - \gamma(1-\delta)\bar{\mu}_1 > [2 + \delta\gamma(q_2 - \frac{1}{2})]\mu_1^B(s_1 = R)$; and
   - $q_2 > \frac{(4+2\delta)\mu_1^B(s_1=R) - 2 + \gamma(1-\delta)}{2\delta}$.

   Note that whenever $1 - \mu_1^B(s_1 = R) > q_2$, the first of the three conditions holds automatically. Note that

$$\begin{aligned}
&\lim_{\delta\to 1} 1 + \tfrac{\gamma}{2}(1+\delta q_2) - \gamma(1-\delta)\bar{\mu}_1 = 1 + \tfrac{\gamma}{2}(1+q_2);\\
&\lim_{\delta\to 1}[2 + \delta\gamma(q_2 - \tfrac{1}{2})]\mu_1^B(s_1 = R) = \left[2 + \gamma(q_2 - \tfrac{1}{2})\right]\mu_1^B(s_1 = R); \text{ and}\\
&\lim_{\delta\to 1} \tfrac{(4+2\delta)\mu_1^B(s_1=R) - 2 + \gamma(1-\delta)}{2\delta} = 3\mu_1^B(s_1 = R) - 1.
\end{aligned}$$

   Since $\mu_1^B(s_1 = R) < \frac{1}{2}$, for sufficiently sufficiently large δ, there exists a positive measure of parameters $p, q_1, q_2, \tilde{\gamma}, \gamma$ $(\tilde{\gamma} \leqslant \gamma)$ for which $q_2 < 1 - \mu_1^B(s_1 = R)$ and the above three conditions hold. The proposition's claim follows. ■.

References

Akerlof, G. A. and Dickens, W. T. (1982). "The Economic Consequences of Cognitive Dissonance", *The American Economic Review*, **72**, 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**, 335-357.

Aronson, E. (1969). "The Theory of Cognitive Dissonance: A Current Perspective", In L. Berkowitz (Ed.), *Advances in Experimental Social Psychology,* **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", *The Quarterly Journal of Economics,* **117(3)**, 871-915.

Blackwell, D. (1950). "Equivalent Comparisons of Experiments", *Annals of Mathematical Statistics*, **21**, 265-272.

Bodner, R. and Prelec, D. (2001). "Self-signaling and Diagnostic Utility in Everyday Decision-making", in Brocas, I. and Carillo, J. D. (Ed.), *Collected Essays in Psychology and Economics* (Oxford: Oxford University Press).

Bracha, A. (2004). "Affective Decision Making In Insurance Markets," mimeo.

Brock, T. C. and Balloun, J. C. (1967). "Behavioral Receptivity to Dissonant Information", *Journal of Personality and Social Psychology*, **6**, 413-428.

Brunnermeier, M. and Parker, J. (2004). "Optimal Expectations," mimeo.

Camerer, C. F. and Weber, R. (1999). "The econometrics and behavioral economics of escalation of commitment: A re-examination of Staw and Hoang's NBA data" *Journal of Economic Behavior and Organization*, **39(1)**, 59-82

Caplin, A. and Leahy, J. (2001). "Psychological Expected Utility Theory and Anticipatory Feelings", *The Quarterly Journal of Economics*, **116**, 55-79.

Carillo, J. and Mariotti, T. (2000). "Strategic Ignorance as a Self-Disciplining Device", *The Review of Economic Studies*, **67**, 529-544.

Compte, O. and Postlewaite, A. (2003), "Confidence and Enhanced Performance", *American Economic Review*, forthcoming.

Constantinides, G. (1990), "Habit Formation: A Resolution of the Equity Premium Bias", *Journal of Political Economy*, **98(3)**, 519-543.

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

Dawson, E., Gilovich, T. and Regan, D. (2004). "Motivated Reasoning and Susceptibility to the "Cell A" Bias", mimeo.

Eliaz, K. and Spiegler, R. (2004). "Can Anticipatory Feelings Explain Anomalous Choices of Information Sources?", mimeo.

Eyster, E. (2003). "Rationalizing the Past: A Taste for Consistency", mimeo.

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**, 74-85.

Fromkin, V. (2000). *Linguistics: An Introduction to Linguistic Theory* (Blackwell).

Fudenberg, D. and Levine, D. K. (1993). "Self-Confirming Equilibrium", *Econometrica*, **61(3)**, 523-545.

Fudenberg, D. and Levine, D. K. (1998). *The Theory of Learning in Games* (MIT Press).

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

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

Graber, D. (1984). *Processing the News: How People Tame the Information Tide* (New York, NY: Longman Press).

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

Karlsson, N., Loewenstein, G., and Seppi, D. (2004). "The 'Ostrich Effect': Selective Exposure to Information About Investments", mimeo.

Koszegi, B. (1999). "Self-image and Economic Behavior", MIT dissertation.

Kruglanski, A. W. (1999). "Motivation, cognition, and reality: Three memos for the next generation of research", *Psychological Inquiry*, **10**, 54-58.

Kunda, Z. (1990). "The Case of Motivated Reasoning", *Psychological Bulletin*, **108**, 480-498.

Larrick, R. P. and Boles, T. L. (1995). "Avoiding Regret in Decisions with Feedback: Negotiation Example", *Organizational Behavior and Human Decision Processes*, **63**, 87-97.

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

Lerman, C., Hughes, C., Trock, B.J., Myers, R.E., Main, D., Bonney, A., Abbaszadegan, M.R., Harty, A.E., Franklin, B.A., Lynch, J.F., and Lynch, H.T. (1999). "Genetic Testing in Families with Hereditary Nonpolyposis Colon Cancer", *Journal of the American Medical Association*, **281**, 1618-1622.

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

McCarthy, A. M., Schoorman, F. D., and Cooper, A. C. (1993). "Reinvestment decisions by entrepreneurs: Rational decision-making or escalation of commitment?", *Journal of Business Venturing*, **8(1)**, 9-24.

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**, 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**, 672-680.

Rabin, M. and Schrag, J. L. (1999). "First Impressions Matter: A Model of Confirmatory Bias", *The Quarterly Journal of Economics*, 37-82.

Schacter, D. L. (1996). *Searching for memory: The brain, the mind, and the past* (New York: Basic Books).

Schoorman, F. D. (1988). "Escalation bias in performance appraisals: An unintended consequence of supervisor participation in hiring decisions", *Journal of Applied Psychology*, **73(1)**, 58-62.

Severin, W. and Tanakard, J. (2000). *Communication Theories: Origins, Methods and Uses in the Mass Media* (Allyn and Bacon).

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**, 27-44.

Staw, B. M., Barsade, S. G., and Koput, K. W. (1997). "Escalation of the credit window: A longitudinal sudy of bank executives' recognition and write-off of problem loans", *Journal of Applied Psychology*, **82(1)**, 130-142.

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**, 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: Academic).

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**, 632-640.

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

Yariv, L. (2002). "I'll See It When I Believe It - A Simple Model of Cognitive Consistency", Cowles Foundation Discussion Paper Number 1352.