

The Dynamics of Chosen Beliefs*

Marc Kaufmann

Mats Köster

November 14, 2024

Abstract

We develop a tractable model of the dynamics of chosen beliefs. An agent derives anticipatory utility from being overly optimistic and overly certain about an underlying state of the world. The agent repeatedly “chooses” what to believe, always naively thinking that this is the one and only time she does so. To feel elated, she must take her chosen belief seriously, both in taking actions and interpreting signals. Our model implies persistent insecurity and fragile optimism in important domains of life as well as dogmatic optimism in less important domains.

Keywords: Belief-Based Utility, Motivated Reasoning, Dogmas, Persistent Insecurity.

*We thank Chiara Aina, Simon Cordes, Katrin Gödker, Paul Heidhues, Botond Kőszegi, Kristóf Madarász, Adam Szeidl, Heidi Thysen, Georg Weizsäcker, and audiences at Berlin, Bonn, and Heidelberg for helpful feedback. Kaufmann: Central European University, KaufmannM@ceu.edu. Köster: Central European University, KosterM@ceu.edu.

1 Introduction

People often distort their beliefs in ways that make them feel better about themselves or the world (Eil and Rao, 2011, Zimmermann, 2020, Drobner, 2022, Möbius et al., 2022). A prominent modeling approach (starting with Brunnermeier and Parker, 2005) formalizes this phenomenon by assuming that at some special point in time, a person resets her belief according to some objective, and from then on carries this new belief. Evidence and introspection suggest, however, that people do so repeatedly. As a case in point, a person who reset her belief exactly once, and were Bayesian otherwise, would generically be unbiased in the long run. We instead take the aforementioned perspective seriously by asking: what if beliefs are reset *repeatedly* according to the same objective?

To address this question, we develop a tractable model of learning for an agent who derives anticipatory utility from beliefs that she can “choose.” The agent wants to be overly optimistic and certain about an underlying state of the world. She repeatedly chooses what to believe, always naively thinking that this is the one and only time she does so. To feel elated the agent must take her chosen belief seriously, both in taking actions and in interpreting signals about the state. We predict a dichotomy of long-run beliefs. An agent who cares only little about the state (in terms of consumption utility) eventually becomes dogmatic and highly overoptimistic. An agent who cares a lot about the state, in contrast, stays forever uncertain, which makes her average optimism fragile.

We introduce our model in Section 2. Consider an agent living for $T \geq 2$ periods. In every period, the agent must guess the state of the world $\theta \sim \mathcal{N}(\mu, \sigma^2)$, which is drawn once and for all. The agent then experiences consumption utility, which depends on the state and her guess, and anticipatory utility from future consumption. The agent incurs a disutility from wrong guesses. But, fixing the utility loss from her guess, the agent’s consumption utility increases in the state.

The agent starts out with a prior *original belief* about the state of the world. In every period, on top of making a guess, the agent can choose her belief. This *chosen belief* shapes how she imagines the future. By choosing to be overly optimistic about the state, the agent can anticipate a brighter future. Moreover, by choosing to be overly certain, she can reduce the anxiety that she derives from possibly making wrong guesses in the future. The agent is constrained, however, in that she has to act on her chosen belief: her chosen belief determines her guess and how she updates upon receiving

a signal on the state. Being overly optimistic, therefore, comes at the cost of making biased guesses. In addition, being overly certain, results in the agent partly ignoring useful information, as we will describe in detail below. Following Brunnermeier and Parker (2005), we assume that the agent uses her original belief in a given period to — consciously or, more likely, subconsciously — trade off the costs and benefits of choosing a different belief.

Our main innovation is allowing for the repeated choice of beliefs, which, as we argue, requires a notion of naivete regarding belief revisions. In-between period t and $t+1$, the agent receives a signal $s_t \in \mathcal{N}(\theta, \nu^2)$ about the state. The agent combines this signal with her chosen belief in period t via Bayes’ rule, which yields her original belief for period $t+1$. This entails the implicit assumption that, upon receiving a signal, the agent forgets that she has chosen her belief in the past. The same assumption is implicit also in existing work. Going beyond the literature, we further assume that, in every period t , the agent naively believes that this is the one and only time she chooses her belief. If the agent anticipated re-choosing her belief in the future, she could hardly feel elated by what she chooses to believe right now. In this sense, our notion of naivete feels psychologically (almost) inevitable, and as we argue in this paper, it has drastic implications for belief dynamics.

We start in Section 3 by analyzing the simplest version of our model with just two periods. This analysis is meant to transparently lay out the mechanics of the model, which will be important for understanding belief dynamics.¹ In particular, we derive a negative relationship between an agent’s level of “optimism” (i.e., her mean belief) and “confidence” (i.e., the variance of her belief): an agent who cares more about the state chooses to be *more optimistic* and *less confident*. The more optimistic the agent chooses to be the more biased her guesses will be. The agent can wash out part of the bias in her second-period guess by placing an excessive weight on the signal. Because the agent updates her chosen belief via Bayes’ rule, however, this requires her to overstate how uncertain she is — she has to “forget” part of what she knows. In the extreme, an agent with a dogmatic prior may even have an incentive to invent uncertainty. This interplay of optimism and confidence has important implications for our main interest: the dynamics of chosen beliefs.

¹ The agent has no incentive to distort her belief in the second, and here last, period because there is no future she can derive anticipatory utility from. In this sense, our model with two periods is a direct application of Brunnermeier and Parker (2005) to a learning problem, and it identifies novel implications of their optimal-expectations framework.

To build some intuition for the dynamics of beliefs, consider Alex, a project-based consultant, who loves to think of herself as being productive. As the highly productive person Alex wants, and therefore believes, to be, she commits herself to taking over major parts of the current project. Overoptimistic Alex soon realizes that it takes her more time than expected to deliver on what she has promised. This humbling experience makes Alex more pessimistic. As time passes however, and the next project is about to start, Alex begins to wonder whether she might have been too harsh with herself in response to just one setback. Because she still likes to think of herself as being productive, Alex chooses to be more optimistic again. At the same time, being worried that she keeps overcommitting to work in the future, Alex also intentionally remains uncertain and thereby (over)responsive to feedback. This raises a natural question: does Alex’s belief ever converge?

We answer this question in Section 4 by studying belief dynamics with an infinite horizon. We compare the long-run chosen belief to three benchmarks, all of which predict that the agent’s long-run belief concentrates on a single point. Our model makes fundamentally different predictions. Depending on the agent’s preferences and the signal structure, she either becomes dogmatic (and increasingly optimistic) or she stays forever uncertain around a “stable” level of overoptimism.

Consider first an agent like Alex who cares a lot about the state. Given her tendency to be optimistic, such an agent is — consciously or, more likely, subconsciously — constantly worried about making biased guesses in the future. As a result, even after observing infinitely many signals, she wants to react to every new signal to correct for this bias going forward. This requires, however, that she stays forever uncertain. We show that the agent’s chosen belief converges to a normally distributed limit belief with an overoptimistic mean and a variance that is bounded away from zero.

At the other extreme, consider an agent who has no preferences over the state. This could be, for example, an impartial scientist who wants to figure out the truth (on, say, polarization) and makes public predictions (on, say, election outcomes) that affect her reputation. Because it reduces her anxiety from making possibly wrong guesses in the future, such an agent has an incentive to be overconfident. As the agent becomes genuinely more certain over time, she eventually convinces herself that she knows the state exactly — she becomes dogmatic and remains biased toward her prior. Here, our model predicts a form of “confirmation bias” (see Benjamin, 2019, for a survey).

Finally, consider the intermediate case of an agent who cares just a little bit about the state. Since such an agent has a moderate incentive to be optimistic, she is less worried (than Alex) about making biased guesses in the future, and thus sees less of a need to remain responsive to signals. Instead, to reduce anxiety, she eventually becomes dogmatic. And once she naively believes to know truth exactly, she chooses to be a little bit more optimistic every period. Because a dogmatic agent ignores the signals, her belief then increases deterministically over time.

We discuss the implications of our results at the hand of comparative statics that distinguish our model from a Bayesian benchmark. A Bayesian’s long-run belief is independent of the incentives she faces, the precision of signals, and her prior belief. All three features of the environment do affect long-run chosen beliefs, however. For example, as we have hinted at above, our agent remains the more uncertain in the long run the more she cares about the state. This persistent insecurity in important domains of life makes the agent’s average optimism fragile, and it generates dramatic — objectively unjustified — swings in beliefs. In particular, irrespective of how overoptimistic the agent has become, a single bad signal (e.g., an overly critical comment by a colleague) can turn her pessimistic. This combination of predictions is consistent with evidence on fragile self-esteem and imposter syndrome (e.g., Ferrari and Thompson, 2006, Berglas, 2006).²

In Section 5, we focus on another central theme in work on motivated reasoning: information avoidance (see Golman et al., 2017, for a survey). As noted by Spiegler (2008), to generate information avoidance in a model of chosen beliefs, we would have to add ad-hoc assumptions on information limiting the ability to choose beliefs (or the anticipatory utility thereof). More fundamentally, information avoidance observed in important settings — such as not getting tested for a genetic disease (Oster et al., 2013) — often seems too costly for being the result of a “rational” trade-off between anticipatory and consumption utility. Instead, we suggest that these findings may be better understood as delaying rather than avoiding information. We show that a naive agent who cares a lot about the state may eternally delay extra information, always planning to get it the next period. Intuitively, an agent who remains forever uncertain prefers to receive information in the future, as otherwise, by overstating her uncertainty, she will “throw away” part of it.

² Alex, for example, is the type of person McKinsey & Company call an “insecure overachiever” (Berglas, 2006).

We discuss our modeling assumptions in Section 6 before we conclude in Section 7 by commenting on the “stability” of long-run beliefs. On the one hand, a dogmatic agent (with preferences over the state) becomes increasingly optimistic over time, and as a result, the signals she observes become harder to reconcile with her original belief. It thus seems plausible that a dogmatic agent eventually “wakes up” and questions her original belief. But, since dogmas are limited to domains that the agent cares about little, her belief diverges slowly and dogmas may stick for a while. Persistent insecurity, on the other hand, seems even more robust to introspection. Because an insecure agent remains forever uncertain, she is less surprised by the signals she observes. This suggests that insecure people may never come to the point of questioning their fragile optimism.

Related literature. Our paper belongs to the theoretical literature on belief-based utility and motivated reasoning (e.g., Akerlof and Dickens, 1982, Caplin and Leahy, 2001, Bénabou and Tirole, 2002, Brunnermeier and Parker, 2005, Kőszegi, 2006, Gottlieb, 2014, Mayraz, 2019, Caplin and Leahy, 2019). We directly build on the model of optimal expectations by Brunnermeier and Parker (2005) as well as contemporaneous work on how people choose beliefs at the cost of worse actions (Landier, 2000, Yariv, 2005). We deviate from existing work in that we study the repeated choice of beliefs. Brunnermeier and Parker (2005) assume that people choose a belief exactly once and are Bayesian otherwise. This assumption is inconsistent with experimental evidence, unless the day of the experiment is the one special day in a subject’s life when she chooses her belief. It also implies that any non-dogmatic belief distortion vanishes in the long run (see Remark 2 or Gottlieb, 2014). Our model, in contrast, predicts persistent belief distortions in the presence of feedback.

On a technical note, our assumption of naivete is similar to the notion of a “naive optimist” in Yariv (2005). Our model differs from Yariv (2005) in two respects, however. First, we assume that the agent chooses her belief to generate anticipatory utility whereas Yariv (2005) models an agent who chooses a belief that makes her feel better about choices taken in the past. Second and more importantly, Yariv (2005) considers a binary state of the world, so that optimism (related to the mean) and confidence (related to the variance) are tied together. This rules out the belief dynamics — based on an interplay of optimism and confidence — that are central to our paper.³

³ Our naivete assumption implies time-inconsistent beliefs. This relates our paper to Brunnermeier et al. (2017) who relax the assumption that optimal expectations must satisfy the law of iterated expectations. Still, Brunnermeier

2 Model

In this section, we simply describe our model. A detailed discussion of our modeling assumptions, equivalent re-formulations, and a comparison with existing models are provided in Section 6.

Utility and beliefs Consider an agent who lives for $T \geq 2$ periods. In each period t , the agent derives *consumption utility* $u(\theta, a_t)$, which depends on her action $a_t \in \mathbb{R}$ and an underlying state of the world $\theta \in \mathbb{R}$. From the perspective of period t , the agent's discounted consumption utility is

$$U(\theta, \mathbf{a}_t) = u(\theta, a_t) + \sum_{\tau=t+1}^T \delta_{\tau-t} u(\theta, a_\tau),$$

where the vector $\mathbf{a}_t = (a_t, a_{t+1}, \dots, a_T)$ collects all relevant (future) actions and $\delta_{\tau-t} \in [0, 1]$ denotes the discount factor that the agent applies to consumption utility τ periods into the future.

Coming into period t , the agent believes that $\theta \sim \mathcal{N}(\mu_t, \sigma_t^2)$ — we call this her *original belief* in period t . The expected discounted consumption utility implied by the agent's original belief is

$$\mathbb{E}_{\mu_t, \sigma_t^2} \left[u(\theta, a_t) + \sum_{\tau=t+1}^T \delta_{\tau-t} u(\theta, a_\tau) \right].$$

To keep the model tractable, we assume that consumption utility takes a linear-quadratic form:

$$u(\theta, a) = \alpha\theta - \frac{1}{2}(a - \theta)^2 \quad \text{for some } \alpha \geq 0. \quad (1)$$

Furthermore, in every period $t \leq T-1$, the agent derives *anticipatory utility* from imagining the future. Following Brunnermeier and Parker (2005), we assume that the agent can manipulate her anticipatory utility by choosing a belief about the state (from the set of all normal distributions). Given a *chosen belief* $\theta \sim \mathcal{N}(\hat{\mu}_t, \hat{\sigma}_t^2)$, the agent's anticipatory utility in period t is

$$\gamma \mathbb{E}_{\hat{\mu}_t, \hat{\sigma}_t^2} \left[\sum_{\tau=t+1}^T \phi_{\tau-t} u(\theta, a_\tau) \right],$$

where $\gamma \in \mathbb{R}_{>0}$ measures the relative importance of anticipatory utility, and $\phi_{\tau-t} \in \mathbb{R}_{\geq 0}$ denotes the weight that the agent puts on imagined consumption utility τ periods into the future.

et al. (2017) maintain the assumption that beliefs are chosen exactly once in order to maximize overall well-being.

Learning environment and updating We set up the simplest possible learning environment. The state θ is drawn once and for all. The agent starts out with some prior belief $\theta \sim \mathcal{N}(\mu_1, \sigma_1^2)$. In-between every two periods, t and $t + 1$, the agent receives an unbiased signal $s_t \in \mathcal{N}(\theta, \nu^2)$.

At the start of period t , the agent chooses a belief $\theta \sim \mathcal{N}(\hat{\mu}_t, \hat{\sigma}_t^2)$, as we describe in detail below. Our key assumption is that the agent has to act on her chosen belief in updating about the state. Specifically, upon observing the signal s_t , the agent updates via Bayes’ rule, using her chosen belief $\mathcal{N}(\hat{\mu}_t, \hat{\sigma}_t^2)$ as the “prior.” This determines the agent’s original belief at the start of period $t + 1$:

$$\mu_{t+1} = (1 - \hat{\lambda}_t)\hat{\mu}_t + \hat{\lambda}_t s_t \quad \text{and} \quad \sigma_{t+1}^2 = \hat{\lambda}_t \nu^2 \quad \text{with} \quad \hat{\lambda}_t := \frac{\hat{\sigma}_t^2}{\hat{\sigma}_t^2 + \nu^2}. \quad (2)$$

The updating we assume differs in two ways from that of a fully Bayesian agent with the same original belief $\mathcal{N}(\mu_t, \sigma_t^2)$. First, in updating her mean belief, our agent starts from a different prior — namely, $\hat{\mu}_t$ instead of μ_t . This entails the implicit assumption that the agent is “naive” in not correcting for the fact that she has chosen her prior. Second, our agent may be more or less certain than her fully Bayesian counterpart and thus places a different weight on the signal — namely, $\hat{\lambda}_t$ instead of $\lambda_t := \frac{\sigma_t^2}{\sigma_t^2 + \nu^2}$. Based on this second distinction, we classify three types of agents.

Definition 1 (Responses to News).

We say that the agent is *dogmatic* if $\hat{\lambda}_t = 0$, *narrow-minded* if $\hat{\lambda}_t \in (0, \lambda_t)$, and *erratic* if $\hat{\lambda}_t \in (\lambda_t, 1]$.

Throughout the paper, we will refer to a bias in the agent’s mean as *over-* and *underoptimism*, respectively, and we will refer to a bias in her variance as either *over-* or *underconfidence*.

Action choice Following the literature, we assume that the agent must pick an action that is optimal under her chosen belief (see Camerer and Lovo, 1999, Allcott et al., 2020, for evidence supporting the idea that people act on their unrealistic beliefs). Formally, we require that

$$a_t = \arg \max_{a \in \mathbb{R}} \mathbb{E}_{\hat{\mu}_t, \hat{\sigma}_t^2} [u(\theta, a)] \quad \text{or, using Eq. (1),} \quad a_t = \hat{\mu}_t. \quad (3)$$

Belief choice Given Eq. (1), the agent can increase anticipatory utility by choosing to be optimistic (high $\hat{\mu}_t$) and certain (low $\hat{\sigma}_t^2$) about the state. Deviating from her original belief, however, comes at the cost of worse inferences from the signals (by Eq. (2)) and worse actions

(by Eq. (3)). Again following Brunnermeier and Parker (2005), we assume that the agent uses her original belief in a period to — consciously or, more likely, subconsciously — trade off these costs (in terms of consumption utility) and benefits (in terms of anticipatory utility).

In contrast to Brunnermeier and Parker (2005), we assume that the agent can choose her belief not only once, but can re-choose it *every* period. In every period t , however, the agent naively believes that this is the one and only time she chooses her belief. As a result, the agent generically mispredicts her future beliefs and actions. Let $\mathcal{N}(\tilde{\mu}_\tau(\mathbf{s}_t^\tau), \tilde{\sigma}_\tau^2(\mathbf{s}_t^\tau))$ be the belief that the agent expects (from the perspective of period t) to hold in period τ for a given sequence of signal realizations $\mathbf{s}_t^\tau := (s_t, \dots, s_{\tau-1})$.⁴ The agent then chooses $\hat{\mu}_t$ and $\hat{\sigma}_t^2$ as to maximize

$$\mathbb{E}_{\hat{\mu}_t, \hat{\sigma}_t^2} \left[u(\theta, \hat{\mu}_t) + \sum_{\tau=t+1}^T \delta_{\tau-t} u(\theta, \tilde{\mu}_\tau(\mathbf{s}_t^\tau)) \right] + \gamma \mathbb{E}_{\hat{\mu}_t, \hat{\sigma}_t^2} \left[\sum_{\tau=t+1}^T \phi_{\tau-t} u(\theta, \tilde{\mu}_\tau(\mathbf{s}_t^\tau)) \right].$$

Parameter restrictions As is standard, we assume $\phi_\tau \geq \phi_{\tau+1}$ and $\delta_\tau \geq \delta_{\tau+1}$ for all $\tau \geq 1$, and we impose $\phi_1, \delta_1 > 0$. Moreover, we bound the total weight on anticipatory utility; formally, $\Phi := \sum_{\tau=1}^{\infty} \phi_\tau < \infty$. We impose similar restrictions on how the agent discounts future consumption utility; in particular, we assume that $\Delta := \sum_{\tau=1}^{\infty} \delta_\tau < \infty$ and $\Omega := \sum_{\tau=1}^{\infty} \delta_\tau \tau < \infty$. This permits, for example, models of exponential and quasi-hyperbolic discounting. Finally, to simplify the exposition of our results, we exclude the knife-edge case in which $\alpha^2 = \frac{\nu^2(1+\Delta)^2}{2\gamma\Phi\Omega}$.

3 Short-Run Beliefs

We start by studying “short-run beliefs,” in the simplest version of our model with two periods. Figure 1 summarizes the timing of events. In the spirit of Spiegel (2019), this section is meant to transparently lay out the mechanics of the model before we move on to our main interest: “long-run beliefs.” To simplify the exposition of our results, throughout this section, we assume $\gamma \leq 2$.

We solve the agent’s problem backwards, starting in the second period. Because the world ends after the second period, the agent does not feel anticipatory utility in $t = 2$. Deviating from her original belief in $t = 2$ has therefore no benefit. Any such deviation still comes at the cost of taking

⁴ In the Online Appendix, we provide more details on the implied updating and the objective function.




Figure 1: *Timing of events with two periods.*

a (perceivedly) worse action a_2 . The agent thus chooses her original belief in $t = 2$. As a result:

Remark 1. If $T = 2$, the agent correctly predicts her second-period belief and action.

Furthermore, because the agent does not feel anticipatory utility in $t = 2$, Remark 1 also implies that our model with two periods coincides with the original model of optimal expectations developed by Brunnermeier and Parker (2005). We can, therefore, think of the results that we derive in this section as identifying novel implications of Brunnermeier and Parker’s model for learning.

We now move on to study the chosen belief in $t = 1$. To ease notation, and without loss, we normalize $\delta_1 = 1$ and $\phi_1 = 1$, and we drop all subscripts referring to the period. Moreover, because there is a one-to-one mapping between the agent’s chosen variance $\hat{\sigma}^2$ and the weight on the signal $\hat{\lambda}$ (see Eq. (2)), it will be convenient to characterize the chosen belief in terms of $\hat{\mu}$ and $\hat{\lambda}$.

Given a chosen belief $\hat{\mu}$ and $\hat{\lambda}$ and a signal realization s , the agent will choose $\tilde{a}(s) = (1 - \hat{\lambda})\hat{\mu} + \hat{\lambda}s$ in $t = 2$. Anticipating her second-period action, the agent chooses $\hat{\mu}$ and $\hat{\lambda}$ as to maximize

$$\underbrace{\alpha\mu - \frac{1}{2}\mathbb{E}_{\mu, \sigma^2}[(\hat{\mu} - \theta)^2]}_{\text{expected consumption utility}} + \underbrace{\alpha\mu - \frac{1}{2}\mathbb{E}_{\mu, \sigma^2}[(1 - \hat{\lambda})\hat{\mu} + \hat{\lambda}s - \theta]^2}_{\text{anticipatory utility}} + \gamma\left(\alpha\hat{\mu} - \frac{1}{2}\hat{\lambda}\nu^2\right). \quad (4)$$

For any given $\hat{\lambda} \in [0, 1]$, the objective in Eq. (4) is strictly concave in $\hat{\mu}$, so the optimal $\hat{\mu}$ solves

$$-(\hat{\mu} - \mu) - (1 - \hat{\lambda})^2(\hat{\mu} - \mu) + \gamma\alpha = 0 \quad \text{or, equivalently,} \quad \hat{\mu} - \mu = \frac{\alpha\gamma}{1 + (1 - \hat{\lambda})^2} =: b(\hat{\lambda}).$$

Proposition 1 (Overoptimism). *For any $\hat{\lambda} \in [0, 1]$, the optimal mean belief equals $\hat{\mu} = \mu + b(\hat{\lambda})$.*

Because consumption utility increases in the state, the agent generates anticipatory utility by being, on average, overly optimistic about the state. Fixing the weight $\hat{\lambda}$, the agent chooses to be more optimistic the more additional consumption utility she derives from an increase in the state (i.e., the higher α) and the more weight she places on anticipatory utility (i.e., the higher γ). At the

same time, the agent understands that overoptimism biases her action not only in the first, but also in the second period. By overresponding to the signal (i.e., by increasing $\hat{\lambda}$), she can counteract her overoptimism and improve her second-period action. The agent’s optimal overoptimism, therefore, increases in the weight she places on the signal, which in turn increases in her chosen uncertainty.

As we show in the Online Appendix, given Proposition 1, the optimal $\hat{\lambda} \in [0, 1]$ maximizes

$$\underbrace{\alpha\gamma b(\hat{\lambda})}_{\text{elation}} - \underbrace{\frac{\gamma}{2}\hat{\lambda}\nu^2}_{\text{anxiety}} - \underbrace{\frac{1}{2}(1 + (1 - \hat{\lambda})^2)b(\hat{\lambda})^2}_{\text{biased actions}} - \underbrace{\frac{1}{2}(\sigma^2 + \nu^2)(\lambda - \hat{\lambda})^2}_{\text{biased learning}}. \quad (5)$$

The first two terms in Eq. (5) reflect how the agent’s choice of certainty affects anticipatory utility. On the one hand, the agent feels elated by imagining a better future, and by Proposition 1, her optimism increases with her uncertainty. This gives an incentive to be overly uncertain. On the other hand, the agent feels anxious about taking a bad second-period action, and convincing herself that she is certain about the state of the world alleviates this anxiety. This implies an incentive to be overly certain. The last two terms in Eq. (5) reflect how the choice of certainty shapes the loss in consumption utility due to biased beliefs and, thus, biased actions. Being overly optimistic directly biases first- and second-period actions. On top, when choosing to be overly (un)certain (i.e., $\hat{\lambda} \neq \lambda$), the agent “misinterprets” the signal, biasing her learning. Biased learning is the more costly the larger is the agent’s overall uncertainty about the state (according to her original belief).

Proposition 2 (Over- and Underconfidence).

- I. The agent chooses to be dogmatic (i.e., $\hat{\lambda} = 0$) if and only if $\alpha \leq \sqrt{\frac{2}{\gamma}(\nu^2 - 2\frac{\sigma^2}{\gamma})}$.
- II. The agent chooses to be erratic (i.e., $\hat{\lambda} > \lambda$) if and only if $\alpha > \sqrt{\frac{\nu^2 + \sigma^2}{2\gamma}(1 + (\frac{\nu^2}{\nu^2 + \sigma^2})^2)}$.
- III. Otherwise the agent chooses to be (weakly) narrow-minded (i.e., $0 < \hat{\lambda} \leq \lambda$).

Depending on her preferences, her prior uncertainty, and the precision of the signal, the agent chooses to be either dogmatic, narrow-minded, or erratic. If the agent cares only little about the state (i.e., $\alpha \approx 0$), she will be either narrow-minded or dogmatic. Such an agent has little incentive to be overoptimistic. Hence, she can be overly certain — alleviating her anxiety — without biasing her second-period action by too much. If such an agent is sufficiently certain already (i.e., σ^2 is low), she fully alleviates her anxiety by becoming dogmatic. If the agent cares a lot about the

state (i.e., α is large), however, she will be erratic. Intuitively, such an agent wants to be highly overoptimistic, and to improve her second-period action, she puts an excessive weight on the signal. Figure 2 illustrates how the chosen uncertainty varies with preferences and prior uncertainty.

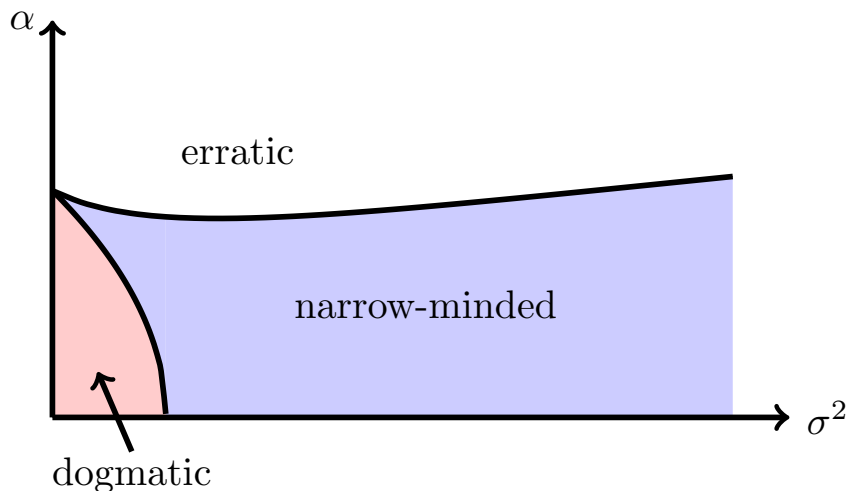


Figure 2: *Optimal (un)certainty depending α and σ^2 for $\gamma^2 = 1$ and $\nu^2 = 2$.*

We conclude our analysis of short-run beliefs with an important remark on erratic updating. By *overresponding* to the signal, an erratic agent effectively “throws away” part of her knowledge. In the extreme, an erratic agent with a dogmatic prior even “makes up” uncertainty.

Corollary 1. *Let $\sigma^2 = 0$. The agent pretends to be uncertain (i.e., $\hat{\sigma}^2 > 0$) if and only if $\alpha > \sqrt{2\frac{\nu^2}{\gamma}}$.*

The logic of Corollary 1 will have important implications for long-run beliefs. An agent who cares enough about the state does *not* want to be certain. This, in turn, suggests that even upon observing infinitely many signals such an agent will remain uncertain — simply because she wants to.

4 Long-Run Beliefs

Our main interest lies in the long-run dynamics of chosen beliefs. For that, we consider our model with an infinite horizon (i.e., $T = \infty$). We start by deriving some benchmark beliefs (Section 4.1); then state our main result (Section 4.2); and conclude with comparative statics (Section 4.3).

4.1 Benchmarks

As a point of comparison, we derive the long-run beliefs implied by three benchmark models. All three benchmarks will have in common that the long-run belief concentrates on a single point.

The first benchmark is that of a correctly specified Bayesian agent. This benchmark corresponds to the limit of our model as $\gamma \rightarrow 0$. The long-run belief converges to the true state θ almost surely.

The second benchmark is that of “misdirected” Bayesian learning. We conceptualize misspecifications in our setting (with exogenous signals) by assuming that the agent believes $s \sim \mathcal{N}(\theta + \xi, \nu^2)$ for some $\xi \in \mathbb{R}$. Such a misspecified Bayesian agent converges to a point belief of $\theta - \xi$.⁵

Our last benchmark is Brunnermeier and Parker’s (2005) model of optimal expectations. Because the agent chooses her belief exactly once, and afterwards updates according to Bayes’ rule, she differs from her fully Bayesian counterpart only in that she starts out with a different prior. By implication, as long as the agent chooses a non-dogmatic prior, she learns the truth eventually.

Remark 2 (Benchmark Beliefs).

- I. The long-run belief of a correctly specified Bayesian agent assigns probability 1 to θ .
- II. The long-run belief of a misspecified Bayesian agent assigns probability 1 to $\theta - \xi$.
- III. Long-run optimal expectations assign probability 1 to either μ_1 or θ .

4.2 Main Result: Long-Run Chosen Beliefs

The long-run belief implied by our model fundamentally differs from the three benchmarks above. Our main result shows that, despite the state being fixed, chosen beliefs do not necessarily converge and even if they do, the limit belief can have full support, reflecting substantial uncertainty.

Formally, the agent’s long-run belief is characterized by $\hat{\mu}_\infty := \lim_{t \rightarrow \infty} \hat{\mu}_t$ and $\hat{\lambda}_\infty := \lim_{t \rightarrow \infty} \hat{\lambda}_t$. In the following, we will distinguish two types of agents, depending on their long-run uncertainty.

Definition 2 (Long-Run Uncertainty).

The agent *develops a dogma* if $\hat{\lambda}_t = 0$ for t large enough; and she *stays forever uncertain* if $\hat{\lambda}_\infty > 0$.

⁵ In general, the belief of a misspecified Bayesian agent concentrates on the outcome distributions that approximately maximize the likelihood of observed signals (e.g., Berk, 1966, Heidhues et al., 2021, Fudenberg et al., 2023).

As a starting point, we verify that in every period $t \in \mathbb{N}$ an optimal belief exists.

Lemma 1 (Existence).

In every period t , an optimal belief exists. The agent chooses $\hat{\lambda}_t \in [0, 1)$ and $\hat{\mu}_t = \mu_t + b(\hat{\lambda}_t)$ with

$$b(\hat{\lambda}_t) := \frac{\alpha\gamma\Phi}{1 + \sum_{\tau=1}^{\infty} \delta_{\tau} \left(\frac{1-\hat{\lambda}_t}{1+(\tau-1)\hat{\lambda}_t} \right)^2}.$$

As in the two-period model analyzed in Section 3, in every period t , the agent chooses to be optimistic, and the bias she introduces in her mean belief increases in how uncertain she chooses to be. Other than in the two-period model, however, the agent's naivete matters. In every period t , the agent believes that this is the one and only time she chooses her belief and, consequently, treats her original belief in period t as "truth." While the agent anticipates that being optimistic today biases her actions in *all* future periods, her naivete causes her to underestimate the (average) bias in future actions. In fact, because the agent re-chooses her belief in every period (on average, becoming more optimistic period by period), the bias in her mean belief accumulates over time.

Lemma 2 (Accumulation of Bias).

The agent's mean belief in period t is normally distributed with an expected value of

$$\mathbb{E}[\hat{\mu}_t] = \mu_1 \prod_{\tau=1}^{t-1} (1 - \hat{\lambda}_{\tau}) + \theta \sum_{\tau=1}^{t-1} \hat{\lambda}_{\tau} \prod_{\ell=\tau+1}^{t-1} (1 - \hat{\lambda}_{\ell}) + \sum_{\tau=1}^t b(\hat{\lambda}_{\tau}) \prod_{\ell=\tau}^{t-1} (1 - \hat{\lambda}_{\ell})$$

and a variance of

$$\text{Var}(\hat{\mu}_t) = \nu^2 \sum_{\tau=1}^{t-1} \hat{\lambda}_{\tau}^2 \prod_{\ell=\tau+1}^{t-1} (1 - \hat{\lambda}_{\ell})^2.$$

With the above, characterizing the agent's long-run belief boils down to understanding how her certainty changes over time. For that, we derive the limit weight that she places on the next signal.

Lemma 3 (Subjective (Un)Learning).

Fix any $\sigma_1^2 \in \mathbb{R}_{\geq 0}$, $\nu^2 \in \mathbb{R}_{> 0}$, $\alpha \in \mathbb{R}_{\geq 0}$, and $\gamma \in \mathbb{R}_{> 0}$. Then, the following statements hold:

- I. The sequence $(\hat{\lambda}_t)_{t \in \mathbb{N}}$ is monotone, and it converges to a limit weight $\hat{\lambda}_{\infty} \in [0, 1)$.*
- II. If $\nu^2 \leq \sigma_1^2$, the sequence $(\hat{\lambda}_t)_{t \in \mathbb{N}}$ is decreasing.*
- III. If $\nu^2 > \sigma_1^2$, there exists some $\underline{\alpha} > 0$, so that $(\hat{\lambda}_t)_{t \in \mathbb{N}}$ is decreasing if and only if $\alpha < \underline{\alpha}$.*

By Part I, for any given signal structure, the agent's certainty monotonically changes over time. Every period, the agent faces the exact same optimization problem, except for her original belief changing over time. Intuitively, the agent chooses to be more uncertain the more uncertain she actually is; that is, the agent's chosen variance increases in the variance of her original belief. It then immediately follows that the sequence of chosen variances and, consequently, the sequence of weights on the next signal are monotone. Because $\hat{\lambda}_t \in [0, 1]$, the sequence $(\hat{\lambda}_t)_{t \in \mathbb{N}}$ is not only monotone, but also bounded, and hence it must converge. Moreover, by Lemma 1, the mean bias introduced in period t is bounded from above by $\alpha\gamma\Phi$. Hence, for any given set of parameters, the agent's chosen mean $\hat{\mu}_t$ does contain some information on the original mean μ_t . As a result, throwing away all existing information by choosing $\hat{\lambda}_t = 1$ is (perceived as) suboptimal, so $\hat{\lambda}_\infty < 1$.

Because the agent chooses her uncertainty to manage the (perceived) bias in her future actions, unlike a Bayesian, she does not necessarily become more certain over time. This depends on the precision of the signals and her preferences. If the signals are precise enough, the variance of the agent's original belief necessarily decreases from σ_1^2 to $\sigma_2^2 = \hat{\lambda}_1\nu^2 \leq \nu^2$ upon observing the first signal. Hence, because she is genuinely less uncertain in the second period compared to the first, the agent also chooses to be less uncertain. Part II then immediately follows from Part I. If the signals are less precise, however, an erratic agent can become more uncertain over time (Part III). As before, the agent chooses to be erratic in order to wash out the bias in future actions, which is increasing in α . Hence, she becomes more uncertain over time if and only if α is large enough.

We can now characterize the agent's long-run belief, and we do so in two separate steps. In a first step, we observe that the agent either develops a dogma or stays forever uncertain.

Proposition 3 (Long-Run Uncertainty).

For any $\sigma_1 \in \mathbb{R}_{\geq 0}$, $\nu^2 \in \mathbb{R}_{> 0}$, and $\gamma \in \mathbb{R}_{> 0}$, there exists some $\bar{\alpha} \in \mathbb{R}_{> 0}$, such that the agent develops a dogma if $\alpha \leq \bar{\alpha}$ or stays forever uncertain otherwise.

To feel less anxious about taking bad actions in the future, an agent who derives little utility from the state being high chooses to be overly certain. And once such an agent is sufficiently certain, she simply convinces herself that she knows the state exactly. On the other hand, an agent who derives a lot of consumption utility from the state being high chooses to be highly optimistic and,

therefore, remains responsive to signals. Building on the logic of Corollary 1, such an agent would never want to be (close to) certain, staying forever uncertain as a result.

In a second step, we derive the distribution of the agent’s long-run mean belief. We identify three types of agents, two of which become dogmatic while the third one stays forever uncertain.

Proposition 4 (Long-Run Mean Belief).

I. If $\alpha = 0$, then $\hat{\mu}_\infty$ is normally distributed with

$$\mathbb{E}[\hat{\mu}_\infty] = (1 - \omega_E)\mu_1 + \omega_E\theta \quad \text{and} \quad \text{Var}(\hat{\mu}_\infty) = \omega_V^2\nu^2 \quad \text{for some} \quad \omega_E, \omega_V \in [0, 1].$$

II. If $\alpha \in (0, \bar{\alpha})$, then $\mathbb{E}[\hat{\mu}_\infty] = \infty$; in particular, for any t large enough,

$$\hat{\mu}_{t+1} - \hat{\mu}_t = \frac{\alpha\gamma\Phi}{1 + \Delta}.$$

III. If $\alpha > \bar{\alpha}$, then $\hat{\mu}_\infty$ is normally distributed with

$$\mathbb{E}[\hat{\mu}_\infty] = \theta + \frac{b(\hat{\lambda}_\infty)}{\hat{\lambda}_\infty} \quad \text{and} \quad \text{Var}(\hat{\mu}_\infty) = \frac{\hat{\lambda}_\infty}{2 - \hat{\lambda}_\infty}\nu^2.$$

An agent without preferences over the state (i.e., $\alpha = 0$) has no incentive to distort her mean belief upward, and as a result, she eventually becomes dogmatic. After a finite number of periods, such an agent holds — and sticks to — a point belief that, on average, lies between her prior mean and the true state of the world (Part I), reflecting a form of “confirmation bias.” In contrast to early models of confirmatory bias (e.g., Rabin and Schrag, 1999, Yariv, 2005), however, the bias implied by our model is driven by considerations regarding the future rather than the past.⁶

As long as $\alpha < \bar{\alpha}$, by Proposition 3, also an agent with preferences over the state becomes dogmatic eventually. By Part II, the mean belief of such an agent diverges. To see why, notice that, in every period t large enough, the agent’s original belief assigns probability 1 to some $\mu_t \in \mathbb{R}$. Treating her original belief as truth, the agent sees little cost in choosing to be slightly more optimistic. And since she derives anticipatory utility from being more optimistic, every period the agent biases her mean belief up a bit more. As a result, her mean belief diverges over time.

⁶ Gottlieb (2014) makes a similar point for an agent who can recode “bad” signals as “good” signals. In contrast to our result, his model suggests that confirmation bias is stronger for agents that care more about the state.

Finally, the mean belief of an agent who cares enough about the state to stay forever uncertain, converges to a limit distribution with full support (Part III). On average, such an agent is overoptimistic in the long-run (i.e., $\mathbb{E}[\hat{\mu}_\infty] > \theta$). At the same time, since she remains uncertain, she keeps responding to every new signal, preventing her belief from concentrating on a single point.

4.3 Comparative Statics and Discussion of the Main Result

We center the discussion of our main result around a set of comparative statics that distinguish our model from the benchmarks analyzed in Section 4.1. The long-run belief of a Bayesian agent, in particular, is independent of (a) the incentives she faces (as modeled by α), (b) the precision of signals she observes, and (c) the prior belief she starts out with. Our model, on the contrary, predicts that the agent’s long-run belief depends on all three aspects of the environment.

We start by studying the role of incentives for long-run beliefs. Suppose that, for a reason outside of our model, it becomes more important to the agent that the state is high. For example, she may start a new job with a compensation scheme that is tightly linked to her productivity θ . Our model can capture such a shift in incentives in reduced-form through an increase in α .

Corollary 2 (Incentives and Long-Run Beliefs). *Fix any $\sigma_1^2 \in \mathbb{R}_{\geq 0}$, $\nu^2 \in \mathbb{R}_{> 0}$, and $\gamma \in \mathbb{R}_{> 0}$.*

- I. $\hat{\lambda}_\infty$ monotonically increases in α , with $\hat{\lambda}_\infty = 0$ for any $\alpha < \bar{\alpha}$ and $\lim_{\alpha \rightarrow \infty} \hat{\lambda}_\infty = 1$.*
- II. $\mathbb{E}[\hat{\mu}_\infty]$ is non-monotonic in α , with $\mathbb{E}[\hat{\mu}_\infty] = \infty$ for any $\alpha < \bar{\alpha}$ and $\lim_{\alpha \rightarrow \infty} \mathbb{E}[\hat{\mu}_\infty] = \infty$.*
- III. For any $\alpha > \bar{\alpha}$, any $\theta \in \mathbb{R}$, and any $\hat{\mu} \in \mathbb{R}$, we have $\mathbb{P}[(1 - \hat{\lambda}_\infty)\hat{\mu} + \hat{\lambda}_\infty s_t < \theta] > 0$.*

The more important it is to the agent that the state is high, the more uncertain she remains in the long run (Part I). We predict that this “insecurity” is paired with (extreme) average optimism as well as huge and objectively unjustified swings in beliefs. More precisely, by Part II, the agent’s average mean belief diverges as $\alpha \rightarrow \infty$. At the same time, by Part III, a single signal — e.g., a critical comment by a colleague — can turn the agent pessimistic, no matter how optimistic she has become. This combination of predictions is consistent with evidence on fragile self-esteem and imposter syndrome: while unrealistically positive self-views are common — in particular, in domains people care about a lot, such as intelligence (Benoît et al., 2015, Charness et al., 2018),

beauty (Eil and Rao, 2011) or job performance (Malmendier and Tate, 2005, Hoffman and Burks, 2020) — many people also struggle with (sometimes unjustified) self doubts, resulting in insecurities and temporal phases of pessimism (Ferrari and Thompson, 2006, Berglas, 2006).⁷




Figure 3: *Properties of limit beliefs for $\delta_1 = \phi_1 = 1$, $\delta_t = \phi_t = 0$ for all $t \geq 2$, $\gamma = 1$, and $\sigma^2 = 1$.*

At the same time, our model predicts extreme overoptimism also in seemingly unimportant domains of life. By Part II, for any $\alpha \leq \bar{\alpha}$, the agent’s mean belief diverges and, as illustrated in Figure 3, if $\hat{\lambda}_\infty$ is continuous in α , also $\lim_{\alpha \searrow \bar{\alpha}} \mathbb{E}[\hat{\mu}_\infty] = \infty$. This form of extreme optimism is markedly different, however, in that it (a) takes time (beliefs diverge over time at a rate proportional to α), and (b) goes along with high confidence and (full) ignorance to evidence. Although extreme, this prediction seems consistent with anecdotal evidence. As a concrete example, 47% of US tennis amateurs above the age of 55 firmly believe that they could win a game against a professional tennis player, which is delusional according to experts. And this is despite the fact that dying on this hill has literally no (career) value for an amateur player above the age of 55.⁸

Next, we study the role of the signal structure. The precision of signals affects how certain the agent chooses to be, which in turn affects her long-run mean belief. Fixing the agent’s original belief, she chooses to be more uncertain the more precise the signals are, simply because she wants

⁷ Kőszegi et al. (2022) propose an equilibrium-based model of fragile self-esteem. The mechanism that generates fragility in their paper is completely different, however. Kőszegi et al. (2022) model an agent who bases her self-esteem on a selection of memories that come to mind. The memories that come to mind depend on how optimistic the agent feels to start with, giving rise to a “personal equilibrium.” A multiplicity of such equilibria is what makes self-esteem fragile. In contrast, the fragility in our model is driven by the fact that the agent remains genuinely insecure.

⁸ Going further, 12% of US tennis amateurs claim that they would sacrifice their marriage for a spot at a Grand Slam tournament, and 20% of US tennis players would give up their entire life savings for the same opportunity. See <https://t1p.de/220pc> and <https://www.youtube.com/watch?v=tiEC1r8n60U> for the survey results (both accessed on March 31, 2024); or have a look at Todd Gallagher’s book on “How Andy Roddick Beat Me with a Frying Pan.”

to put more weight on a more precise signal. But, all else equal, observing a more precise signal makes the agent more certain, and as a result, she also chooses to be more certain going forward. While impact of signal precision on the agent's long-run belief is therefore ambiguous in general, our next result identifies the effects of observing highly (im)precise signals.

Corollary 3 (Signal Precision and Long-Run Beliefs). *Fix any $\sigma_1^2 \in \mathbb{R}_{\geq 0}$, $\alpha \in \mathbb{R}_{>0}$, and $\gamma \in \mathbb{R}_{>0}$.*

- I. There exists some $\bar{\nu}^2 > 0$ such that, for any $\nu^2 > \bar{\nu}^2$, $\hat{\lambda}_\infty = 0$ and $\mathbb{E}[\hat{\mu}_\infty] = \infty$.*
- II. $\lim_{\nu^2 \rightarrow 0} \hat{\lambda}_\infty = 1$ and $\lim_{\nu^2 \rightarrow 0} \mathbb{E}[\hat{\mu}_\infty] = \theta + \alpha\gamma\Phi$.*

By Part I, if the signals are sufficiently imprecise, the agent becomes dogmatic eventually, because the informational content of a signal does not make up for the anxiety from remaining uncertain. At the other extreme, as signals become arbitrarily precise, the agent remains maximally uncertain (Part II). Intuitively, in this case the anxiety from being uncertain vanishes: $\hat{\lambda}_t \nu^2 \rightarrow 0$ as $\nu^2 \rightarrow 0$. Hence, by choosing $\hat{\lambda}_t = 1$, the agent can fully debias herself every period at approximately no cost.

Finally, we discuss the role of the agent's prior for her long-run belief.

Corollary 4 (Prior and Long-Run Beliefs).

For any $\nu^2 \in \mathbb{R}_{>0}$, $\alpha \in \mathbb{R}_{\geq 0}$, and $\gamma \in \mathbb{R}_{>0}$, $\hat{\lambda}_\infty$ is independent of μ_1 and weakly increases in σ_1^2 .

Because of the quadratic loss from inaccurate actions, the agent's prior mean has no effect on how uncertain she chooses to be. In combination with Proposition 4, this also implies that, for any $\alpha > 0$, the agent's prior mean has no effect on her long-run mean belief. The agent's prior uncertainty, in contrast, may affect long-run uncertainty and as a result, also the long-run mean belief. In general, the agent's long-run uncertainty weakly increases in her prior uncertainty, and as illustrated in Figure 4, σ_1^2 determines whether the agent remains forever uncertain or not.

We conclude this section by studying in more detail the confirmation bias implied by our model. As an illustration, consider an impartial scientist who wants to figure out the truth (i.e., $\alpha = 0$). Every period she runs a (thought) experiment that generates a signal s_t about the true state of the world, and makes a testable prediction about the state. Consider, for example, a political scientist studying polarization and predicting election outcomes or a climatologist making predictions about changes in sea levels or precipitation patterns. The scientist's predictions are scrutinized by the




Figure 4: *Chosen uncertainty* for $\delta_1 = \phi_1 = 1$, $\delta_t = \phi_t = 0$, $t \geq 2$, $\nu^2 = 25$, $\alpha = 6.8$, and $\gamma = 1$.

scientific community or general public, and she incurs a quadratic loss (in reputation) from making wrong predictions. Even though the scientist is impartial, she will not learn the truth. Instead, because she is anxious to make wrong predictions, the scientist eventually develops a dogma and remains biased toward her initial belief. Formally, the agent’s long-run mean belief increases in her prior mean μ_1 , with this bias in her mean belief being larger the lower is her prior variance σ_1^2 .

5 Information Preferences

A central theme in work on motivated reasoning is preferences for information. Oster et al. (2013), for example, find that people with a genetic predisposition for Huntington’s disease avoid getting tested at the cost of making worse decisions in other domains of life. As already pointed out in Spiegler (2008, 2019), a model of chosen beliefs, like ours, cannot explain “information avoidance” (see Golman et al., 2017, for a survey) in this classical sense. In this section, we first re-state Spiegler’s observation using our notation, and then study preferences over the timing of information, deriving a novel form of “information delay.” Arguably, suboptimal behavior in high-stakes decisions such as genetic testing are better understood as information delay than avoidance.

We start by re-stating Spiegler’s observation that a model like ours generates a (weak) preference

for information. To avoid clutter, consider the version of our model with two periods, and again normalize $\delta_1 = \phi_1 = 1$. We denote by $\mathcal{U}(\sigma_1^2, \nu^2; \alpha, \gamma)$ the agent's (ex-ante) value function, and calculate the marginal value of information. Going beyond Spiegler (2008), we distinguish between preferences for “immediate information,” which we conceptualize as a reduction in prior variance σ_1^2 , and “future information,” which we conceptualize as a reduction in signal variance ν^2 .

Proposition 5 (The Marginal Value of Information).

- I. $\frac{\partial}{\partial \sigma_1^2} \mathcal{U}(\sigma_1^2, \nu^2; \alpha, \gamma) < 0$ and $\frac{\partial}{\partial \nu^2} \mathcal{U}(\sigma_1^2, \nu^2; \alpha, \gamma) \leq 0$ almost everywhere.
- II. $\frac{\partial^2}{\partial \alpha \partial \sigma_1^2} \mathcal{U}(\sigma_1^2, \nu^2; \alpha, \gamma) \geq 0$ and $\frac{\partial^2}{\partial \alpha \partial \nu^2} \mathcal{U}(\sigma_1^2, \nu^2; \alpha, \gamma) \leq 0$ almost everywhere.
- III. For any $\gamma \in \mathbb{R}_{>0}$, there exists an $\check{\alpha} \in \mathbb{R}_{>0}$, such that for all $\alpha > \check{\alpha}$, $\sigma_1^2 \in \mathbb{R}_{\geq 0}$, and $\nu^2 \in \mathbb{R}_{>0}$,

$$\left| \frac{\partial}{\partial \sigma_1^2} \mathcal{U}(\sigma_1^2, \nu^2; \alpha, \gamma) \right| < \left| \frac{\partial}{\partial \nu^2} \mathcal{U}(\sigma_1^2, \nu^2; \alpha, \gamma) \right| \quad (a.e.).$$

By Part I, the agent values both immediate and future information, and would thus never avoid extra information. By Part II, the value of immediate information (weakly) decreases in α . This is because an agent with a higher α places a larger (potentially excessive) weight on the signal and thus throws away a larger part of the information she receives before choosing her belief. For the same reason, the value of future information (weakly) increases in α . Intuitively, an increase in signal precision is more valuable the larger is the weight that the agent places on the signal. By Part III, if α is large enough, the value of future information exceeds that of immediate information.

We use this last observation to show that our model can generate “information delay.” For that, consider again the model with an infinite horizon, and suppose that the agent can once receive an additional (unbiased, normal) signal with precision d/ν^2 for some $d \in \mathbb{N}$. This additional signal could be, for example, a medical test while the “regular” signals are symptoms. In every period t , the agent decides whether to receive the additional signal now — before choosing a new belief — or delay the additional signal to a later period. Notice that a Bayesian agent would never delay the signal, as delaying it would imply that the current decision is worse than it could be. In contrast, as illustrated in Proposition 5, an erratic agent may prefer getting information in the future, because she understands that otherwise she would throw away part of it. In the extreme, an erratic agent may therefore delay the additional signal eternally.

Proposition 6 (Eternally Delaying Information).

Let $T = \infty$, $\delta_1 = \phi_1 = 1$, and $\delta_t = \phi_t = 0$ for all $t \geq 2$. For any $\sigma_1^2 \leq \nu^2$ and any $\gamma \in \mathbb{R}_{>0}$, there exists some $\hat{\alpha} \in \mathbb{R}_{>0}$, such that any agent with $\alpha > \hat{\alpha}$ eternally delays the additional signal.

To understand the implications of Proposition 6, it is useful to contrast our model with Brunnermeier and Parker (2005), according to which the agent can choose her belief exactly once. Also when choosing a belief exactly once, the agent may want to delay information, but she will receive the signal the latest in $t = 2$. In this sense, information delay is not too costly in Brunnermeier and Parker (2005). While our agent also always plans to receive the signal in the next period, she does not follow through with this plan and instead delays the signal over and over again.⁹

6 Discussion of Modeling Assumptions

Belief utility Building on a sizable literature in economics (e.g., Jevons, 1905, Loewenstein, 1987, Caplin and Leahy, 2001, Brunnermeier and Parker, 2005), and an even larger one in psychology, we model an agent who derives *anticipatory utility* from future consumption. We take a particular stand on how the agent imagines the future: she forms a coherent view of the future shaped by her period- t chosen belief, placing a weight of $\gamma\phi_{\tau-t} \geq 0$ on consumption τ periods into the future.

Deviating from Brunnermeier and Parker (2005), we model an agent who does *not* anticipate future anticipatory utility.¹⁰ This is in the tradition of Jevons (1905), who thought of people as maximizing *current* happiness. To understand the implications of this assumption, notice that every signal, on average, drags down the agent’s overoptimistic belief. Hence, anticipated anticipatory utility would be lower than today’s anticipatory utility, and anticipating anticipatory utility means imagining the future using multiple, incoherent beliefs. By instead assuming that anticipatory utility is based on a single, coherent belief, our model is closer to the cover version of “optimal expectations” studied in Spiegel (2008, 2019) and other models of “wishful thinking” like Caplin and Leahy (2019). And, as we show in the Online Appendix, this assumption has bite: assuming

⁹ Notice that Proposition 6 is not driven by the extreme discounting we assumed. As we show in the Online Appendix, with less extreme discounting, delaying information becomes even more attractive because relative to receiving information right now, it improves *all* future actions and thereby allows for a larger bias in the mean belief.

¹⁰ Moreover, we assume that the agent does not feel “memory utility” from remembering the past.

that the agent anticipates anticipatory utility (in combination with our parametric assumptions discussed below) necessarily implies a dogmatic long-run belief.

People may, of course, derive utility from beliefs for reasons other than anticipatory utility (as modeled, for example, in Yariv, 2005, Kőszegi, 2006). For instance, people may want to hold certain beliefs in order to not admit bad decisions in the past. While our framework in principle allows us to study learning with such alternative forms of belief utility, we have not done so yet.

Cost of belief distortions Following Brunnermeier and Parker (2005), we introduce an implicit cost of distorting one’s belief, which comes in the form of worse actions and lower consumption utility. Other models, like Caplin and Leahy (2019), introduce an explicit cost of deviating from one’s original belief. As we show in the Online Appendix, we can re-write our model in terms of such an explicit cost. Our “cost function” is shaped by consumption utility, and given our functional form assumptions, shares some properties with cost functions used in the literature.

Restrictions on consumption utility and chosen beliefs We assume linear-quadratic consumption utility, and allow any normally distributed chosen belief. Given this functional form, the agent has preferences over her mean belief and variance only. Choosing a normally distributed belief is then indeed optimal, and our restriction to a normally distributed beliefs without loss.

It is worth highlighting a couple of advantages our setup has. First, Bayesian updating with a normal prior and normally distributed signals is simple, which buys us tractability. Second, and more importantly, mean and variance of a normal distribution are independent. In contrast to existing work focusing on binary outcomes (Yariv, 2005, Oster et al., 2013), we can thus study how optimism (i.e., distortions of the mean) and confidence (i.e., distortions of the variance) interact.

Finally, we deviate from Brunnermeier and Parker (2005) in that we allow the chosen belief to have a wider support than the agent’s original belief (cf Assumption 1(iv) in their paper). As shown in Corollary 1, their assumption is with loss of generality: the original belief can have a narrower support than the chosen one (namely, $\{\mu\}$ instead of \mathbb{R}). If we instead adopted Brunnermeier and Parker’s assumption, we would effectively assume that a dogmatic agent cannot choose her belief. Under this alternative assumption, the mean belief of a dogmatic agent would not diverge. With

this one exception, however, all of our results continue to hold.¹¹

Naivete We impose three forms of naivete, two of which reflect standard assumptions in the literature. First, we assume that the agent is naive in a “backward-looking” sense: in every period, she treats her original belief as truth, not correcting for the fact that she has (partly) chosen it in the past. Second, and relatedly, we assume that the agent does not make inferences from how her imagined and realized utility differ. By forcing the agent’s beliefs to satisfy the law of iterated expectations (Part (iii) of their Assumption 1), Brunnermeier and Parker (2005) implicitly make the same assumptions. As a natural first step, we maintain both assumptions. Allowing the agent to question her original belief — potentially based on the divergence of imagined and realized utility¹² — raises interesting theoretical issues that seem orthogonal to our main interests.

Third, we assume that the agent is naive in a “forward-looking” sense: in every period, she thinks that this is *the one and only time* she chooses her belief. Intuitively, for the agent to feel elated through the choice of beliefs, she must take her chosen belief seriously, which seems impossible when she already anticipates re-choosing her belief in the future. Brunnermeier and Parker (2005) circumvent this problem by allowing the agent to choose her belief *exactly once*. Hence, by construction, the agent is correctly specified despite taking her chosen belief seriously. By doing so, Brunnermeier and Parker (2005) make the highly unrealistic assumption that there is a single special point in time at which a person chooses her belief. This implies, as we argue in Section 4, that almost any belief distortion à la Brunnermeier and Parker (2005) washes out in the long run. We, in contrast, want to take seriously the idea that people can re-choose their beliefs, and assume forward-looking naivete to make the agent take her chosen beliefs seriously. Yariv (2005) studies the implications of a similar form of naivete in a different setting.

Trigger An important shortcoming of our model is that we lack a theory of what triggers the choice of new beliefs. The current model ties the choice of beliefs to receiving a signal. This assumption is not grounded in psychological evidence, however. Still, as long as the agent keeps

¹¹ The agent — even if aware that becoming dogmatic prevents her from re-choosing her belief in the future — would not choose a different belief because she anyways thinks that this is the one and only time she does so.

¹² Notice, however, that this inference is complicated through anticipatory utility masking low consumption utility.

re-choosing her belief on a regular basis, all our results remain to hold. Formally, suppose that the agent re-chooses her belief every $d + 1 \in \mathbb{N}$ periods, so that she observes d signals in-between. Since our agent updates according to Bayes’ rule, observing d normally distributed signals with precision $1/\nu^2$ is equivalent to observing a single signal with precision d/ν^2 . Hence, a more infrequent, but regular choice of beliefs is equivalent to making the signals more precise, and because we derive results for an arbitrary precision of the signals, our model subsumes this case. In other settings, however, the exact trigger might be important, so the question calls for future research.

Distorting prior vs. signal structure Following Brunnermeier and Parker (2005), we assume that the agent chooses her prior belief and then updates via Bayes’ rule taking the chosen belief and the signal structure as given. This identifies the beliefs that the agent *wants* to hold over time, ignoring potential constraints stemming from the psychological mechanism of forming such beliefs. A natural alternative — as in Rabin and Schrag (1999) or Bénabou and Tirole (2002) — is to assume that the agent takes her prior as given and manipulates her belief about the signal structure. As we show in the Online Appendix, if the agent chooses her perception of the signal structure in order to maximize anticipatory utility (see also Gottlieb, 2014), such an alternative model yields testable predictions that differ from those of our model.¹³

7 Concluding Remarks

We conclude by discussing a question orthogonal to the analysis so far: (when) does the agent “wake up” and question her original belief? A dogmatic agent becomes increasingly optimistic over time, and as a result, the observed signals become increasingly unlikely under her original belief. This suggests that she may question her (extreme) dogmatic belief eventually. At the same time, because dogmas are limited to domains the agent cares about fairly little, dogmatic beliefs get detached from reality slowly. Dogmas may, therefore, stick for a prolonged period of time.

Persistent insecurity seems even more robust to introspection in the sense that the next signal is less unrealistic under the agent’s original belief. First, an insecure agent is less overoptimistic in

¹³ Modeling an agent who optimally chooses her belief about the signal structure upon observing the signal might provide a microfoundation for some quasi Bayesian models used in the literature (e.g., Möbius et al., 2022).

the long run, which makes an unbiased signal less surprising. Second, the agent remains uncertain in the long run, which again makes an unbiased, on average negative, signal less surprising. This suggests that insecure people might never come to the point of questioning their own insecurity.

A formal model of waking up is beyond the scope of the present paper, however. Such a model would require us to take a stand on what the agent does upon “rejecting” her original belief. Existing work (Gagnon-Bartsch et al., 2023, Ba, 2024) has studied waking up in the presence of an exogenously given, natural alternative that is (infinitely) more likely in the light of the data. While plausible in the context of a (misspecified) Bayesian model, it is not obvious how to combine this idea with a model of chosen beliefs. We therefore leave this question to future research.

Appendix: Proofs

Proof of Proposition 1. The proof is given in the text. □

Proof of Proposition 2. Combing the first and third term in (5), the agent’s objective simplifies to

$$\frac{\alpha^2 \gamma^2}{1 + (1 - \hat{\lambda})^2} - (\lambda - \hat{\lambda})^2 (\sigma^2 + \nu^2) - \gamma \hat{\lambda} \nu^2. \quad (6)$$

First, we show that for any $\gamma \leq 2$, the objective function in Eq. (6) is single-peaked in $\hat{\lambda}$. The first derivative of Eq. (6) with respect to $\hat{\lambda}$ is

$$\alpha^2 \gamma^2 \frac{1 - \hat{\lambda}}{(1 + (1 - \hat{\lambda})^2)^2} + (\lambda - \hat{\lambda})(\sigma^2 + \nu^2) - \frac{\gamma}{2} \nu^2, \quad (\text{FD})$$

and its second derivative with respect to $\hat{\lambda}$ is

$$\alpha^2 \gamma^2 \frac{3\hat{\lambda}^2 - 6\hat{\lambda} + 2}{(1 + (1 - \hat{\lambda})^2)^3} - (\sigma^2 + \nu^2). \quad (\text{SD})$$

The second derivative (SD) is decreasing in $\hat{\lambda}$, it equals

$$\frac{\alpha^2 \gamma^2}{4} - (\sigma^2 + \nu^2)$$

at $\hat{\lambda} = 0$, and it is negative for any $\hat{\lambda} > \frac{3-\sqrt{3}}{3} \approx 0.42$. This implies that Eq. (6) is either concave or first convex and then concave. Because the first derivative is negative for sufficiently large $\hat{\lambda}$, it

follows that Eq. (6) is single-peaked, with the maximizer solving the first-order condition, whenever $\text{FD}|_{\hat{\lambda}=0} > 0$. Finally, suppose the first derivative is non-positive at $\hat{\lambda} = 0$, which is equivalent to

$$\frac{\alpha^2 \gamma^2}{4} + \sigma^2 - \frac{\gamma}{2} \nu^2 \leq 0.$$

We then conclude that, for any $\gamma \leq 2 + 4\frac{\sigma^2}{\nu^2}$, $\text{FD}|_{\hat{\lambda}=0} \leq 0$ implies $\text{SD}|_{\hat{\lambda}=0} \leq 0$. Hence, whenever $\text{FD}|_{\hat{\lambda}=0} \leq 0$, Eq. (6) is single-peaked at $\hat{\lambda} = 0$.

We now conclude that, for any $\gamma \leq 2$, the agent is dogmatic (i.e., $\hat{\lambda} = 0$) if and only if

$$\text{FD}|_{\hat{\lambda}=0} \leq 0 \quad \text{or, equivalently,} \quad \gamma \geq 2\frac{\sigma^2}{\nu^2} \quad \text{and} \quad \alpha^2 \leq 2\frac{\nu^2}{\gamma} - 4\frac{\sigma^2}{\gamma^2}.$$

Moreover, the agent is erratic (i.e., $\hat{\lambda} > \lambda$) if and only if

$$\text{FD}|_{\hat{\lambda}=\lambda} > 0 \quad \text{or, equivalently,} \quad \alpha^2 > \frac{1}{\gamma} \frac{\nu^2 (1 + (1 - \lambda)^2)^2}{2(1 - \lambda)}.$$

Otherwise the agent is (weakly) narrow-minded (i.e., $0 < \hat{\lambda} \leq \lambda$). This completes the proof. \square

Proof of Corollary 1. Because $\sigma^2 = 0$ implies $\lambda = 0$, it follows from Part II of Proposition 2. \square

Proof of Lemma 1. Define $w_t^\tau(\hat{\lambda}_t) := \frac{\hat{\lambda}_t(\tau-t)}{1+(\tau-(t+1))\hat{\lambda}_t}$. As we show in the Online Appendix, the agent chooses \hat{b}_t and $\hat{\lambda}_t$ as to maximize

$$\begin{aligned} -\frac{1}{2}\hat{b}_t^2 - \frac{1}{2} \sum_{\tau=t+1}^{\infty} \delta_{\tau-t} \left((w_t^\tau(\lambda_t) - w_t^\tau(\hat{\lambda}_t))^2 \left(\sigma^2 + \frac{\nu^2}{\tau-t} \right) + (1 - w_t^\tau(\hat{\lambda}_t))^2 \hat{b}_t^2 \right) \\ + \gamma \sum_{\tau=t+1}^{\infty} \phi_{\tau-t} \left(\alpha \hat{b}_t - \frac{1}{2} w_t^\tau(\hat{\lambda}_t) \frac{\nu^2}{\tau-t} \right). \end{aligned} \quad (7)$$

We start by solving for the optimal \hat{b}_t as a function of $\hat{\lambda}_t$. Fixing any $\hat{\lambda}_t \in [0, 1]$, the above objective function is strictly concave in \hat{b}_t . Hence, the optimal \hat{b}_t satisfies the first-order condition

$$-\hat{b}_t \left(1 + \sum_{\tau=t+1}^{\infty} \delta_{\tau-t} (1 - w_t^\tau(\hat{\lambda}_t))^2 \right) + \gamma \alpha \underbrace{\sum_{\tau=t+1}^{\infty} \phi_{\tau-t}}_{= \Phi} = 0,$$

which in turn implies

$$\hat{b}_t = \frac{\alpha \gamma \Phi}{1 + \sum_{\tau=t+1}^{\infty} \delta_{\tau-t} (1 - w_t^\tau(\hat{\lambda}_t))^2} = \frac{\alpha \gamma \Phi}{1 + \sum_{\tau=1}^{\infty} \delta_{\tau} \left(\frac{1 - \hat{\lambda}_t}{1 + (\tau-1)\hat{\lambda}_t} \right)^2}. \quad (8)$$

Plugging (8) into (7), we conclude that the agent chooses $\hat{\lambda}_t \in [0, 1]$ as to maximize

$$f(\hat{\lambda}_t, \lambda_t, \alpha) := \frac{1}{2} \frac{\alpha^2 \gamma^2 \Phi^2}{1 + \sum_{\tau=t+1}^{\infty} \delta_{\tau-t} (1 - w_t^\tau(\hat{\lambda}_t))^2} - \frac{\gamma}{2} \sum_{\tau=t+1}^{\infty} \phi_{\tau-t} w_t^\tau(\hat{\lambda}_t) \frac{\nu^2}{\tau-t} \\ - \frac{1}{2} \sum_{\tau=t+1}^{\infty} \delta_{\tau-t} (w_t^\tau(\lambda_t) - w_t^\tau(\hat{\lambda}_t))^2 \left(\sigma_t^2 + \frac{\nu^2}{\tau-t} \right).$$

Notice that

$$\frac{\partial}{\partial \hat{\lambda}_t} w_t^\tau(\hat{\lambda}_t) = \frac{(\tau-t)}{(1 + (\tau - (t+1))\hat{\lambda}_t)^2},$$

and because $\sigma_t^2/\nu^2 = \lambda_t/(1-\lambda_t)$,

$$(w_t^\tau(\lambda_t) - w_t^\tau(\hat{\lambda}_t)) \left(\sigma_t^2 + \frac{\nu^2}{\tau-t} \right) = \frac{(\lambda_t - \hat{\lambda}_t)(\tau-t)}{(1 + (\tau - (t+1))\lambda_t)(1 + (\tau - (t+1))\hat{\lambda}_t)} \frac{\nu^2}{\tau-t} \frac{1 + (\tau - (t+1))\lambda_t}{1 - \lambda_t} \\ = \nu^2 \frac{\lambda_t - \hat{\lambda}_t}{1 - \lambda_t} \frac{1}{1 + (\tau - (t+1))\hat{\lambda}_t}.$$

With this, we can re-write the first partial derivative of f with respect to $\hat{\lambda}_t$ as follows

$$\frac{\partial f}{\partial \hat{\lambda}_t} = \alpha^2 \gamma^2 \Phi^2 \frac{\sum_{\tau=t+1}^{\infty} \delta_{\tau-t} (\tau-t) \frac{(1-\hat{\lambda}_t)}{(1+(\tau-(t+1))\hat{\lambda}_t)^3}}{\left(1 + \sum_{\tau=t+1}^{\infty} \delta_{\tau-t} \left(\frac{1-\hat{\lambda}_t}{1+(\tau-(t+1))\hat{\lambda}_t} \right)^2 \right)^2} - \frac{\gamma}{2} \nu^2 \sum_{\tau=t+1}^{\infty} \phi_{\tau-t} \frac{1}{(1 + (\tau - (t+1))\hat{\lambda}_t)^2} \\ + \nu^2 \sum_{\tau=t+1}^{\infty} \delta_{\tau-t} \frac{(\tau-t)}{(1 + (\tau - (t+1))\hat{\lambda}_t)^3} \frac{\lambda_t - \hat{\lambda}_t}{1 - \lambda_t}.$$

We conclude by making two observations on $\hat{\lambda}_t$. First, because f is continuous in $\hat{\lambda}_t$, and because $\hat{\lambda}_t$ is chosen from the closed interval $[0, 1]$, an optimal $\hat{\lambda}_t$ exists. Second, for any $\lambda_t < 1$, $\frac{\partial f}{\partial \hat{\lambda}_t} \Big|_{\hat{\lambda}_t=1} < 0$. Hence, for any $\lambda_t < 1$, the optimal weight on the next signal satisfies $\hat{\lambda}_t < 1$. \square

Proof of Lemma 2. The result follows immediately when (iteratively) applying the updating rule specified in Section 2 and using the fact that $\mathbb{E}[s_t] = \theta$ and $\text{Var}(s_t) = \nu^2$ for all t . \square

Proof of Lemma 3. (Throughout, we will fix $\nu^2 \in \mathbb{R}_{>0}$. Acknowledging a slight imprecision, when taking a partial derivative with respect to λ_t , we refer to changes in λ_t for a fixed ν^2 .)

Part I. We have to show that *the sequence $(\hat{\lambda}_t)_{t \in \mathbb{N}}$ is monotone*. Using the same notation as in the proof of Lemma 1, we observe that, for any $\hat{\lambda}_t < 1$,

$$\begin{aligned} \frac{\partial^2 f}{\partial \lambda_t \partial \hat{\lambda}_t} &= \nu^2 \sum_{\tau=t+1}^{\infty} \delta_{\tau-t} \frac{(\tau-t)}{(1+(\tau-(t+1))\hat{\lambda}_t)^3} \frac{\partial}{\partial \lambda_t} \left(\frac{\lambda_t - \hat{\lambda}_t}{1 - \lambda_t} \right) \\ &= \nu^2 \sum_{\tau=t+1}^{\infty} \delta_{\tau-t} \frac{(\tau-t)}{(1+(\tau-(t+1))\hat{\lambda}_t)^3} \frac{1 - \hat{\lambda}_t}{(1 - \lambda_t)^2} > 0. \end{aligned}$$

Hence, f satisfies *strictly increasing differences* in $(\hat{\lambda}_t, \lambda_t)$: i.e., for any $\hat{\lambda}_t \geq \hat{\lambda}'_t$ and $\lambda_t \geq \lambda'_t$,

$$f(\hat{\lambda}_t, \lambda_t, \alpha) - f(\hat{\lambda}'_t, \lambda_t, \alpha) \geq f(\hat{\lambda}_t, \lambda'_t, \alpha) - f(\hat{\lambda}'_t, \lambda'_t, \alpha),$$

holding with a strict inequality if $\hat{\lambda}_t > \hat{\lambda}'_t$ and $\lambda_t > \lambda'_t$. By Theorem 10.6 in Sundaram (1996), therefore, $\hat{\lambda}_t$ monotonically increases in λ_t for any $t \geq 1$. By assumption, for any $t \geq 2$, we have $\lambda_t = \frac{\hat{\lambda}_{t-1}}{1 + \hat{\lambda}_{t-1}}$, which increases in $\hat{\lambda}_{t-1}$. Hence, for any $t \geq 2$, $\hat{\lambda}_t$ monotonically increases in $\hat{\lambda}_{t-1}$.

To verify monotonicity, we distinguish two cases. First, suppose that $\lambda_1 > \frac{\hat{\lambda}_1}{1 + \hat{\lambda}_1} = \lambda_2$. Because $\hat{\lambda}_t$ is increasing in λ_t , we have $\hat{\lambda}_2 < \hat{\lambda}_1$. And, because $\hat{\lambda}_t$ monotonically increases in $\hat{\lambda}_{t-1}$, we further conclude that $\hat{\lambda}_t < \hat{\lambda}_{t-1}$ for all $t \geq 3$. Second, suppose that $\lambda_1 \leq \frac{\hat{\lambda}_1}{1 + \hat{\lambda}_1} = \lambda_2$. An argument similar to that in the first case implies $\hat{\lambda}_t \geq \hat{\lambda}_{t-1}$. Hence, the sequence $(\hat{\lambda}_t)_{t \in \mathbb{N}}$ is monotone.

The sequence $(\hat{\lambda}_t)_{t \in \mathbb{N}}$ is not only monotone, but because $\hat{\lambda}_t \in [0, 1]$, it is also bounded. Hence, by the monotone convergence theorem, the sequence $(\hat{\lambda}_t)_{t \in \mathbb{N}}$ converges to some $\hat{\lambda}_\infty \in [0, 1]$. It remains to be shown that $\hat{\lambda}_\infty < 1$. Notice that $\lambda_\infty = \frac{\hat{\lambda}_\infty}{\hat{\lambda}_\infty + 1}$ and that $\hat{\lambda}_\infty$ solves $\frac{\partial f}{\partial \hat{\lambda}_t} \Big|_{\hat{\lambda}_t = \hat{\lambda}_\infty, \lambda_t = \lambda_\infty} = 0$. Now, for the sake of a contradiction, assume that $\hat{\lambda}_\infty = 1$ and thus $\lambda_\infty = \frac{1}{2}$. Then,

$$\frac{\partial f}{\partial \hat{\lambda}_t} \Big|_{\hat{\lambda}_t = \hat{\lambda}_\infty, \lambda_t = \lambda_\infty} = -\frac{\gamma}{2} \nu^2 \sum_{\tau=t+1}^{\infty} \phi_{\tau-t} \frac{1}{(\tau-t)^2} - \nu^2 \sum_{\tau=t+1}^{\infty} \delta_{\tau-t} \frac{1}{(\tau-t)^2} < 0;$$

a contradiction. Hence, $(\hat{\lambda}_t)_{t \in \mathbb{N}}$ converges to some $\hat{\lambda}_\infty \in [0, 1)$.

Part II. By Lemma 1, we have $\hat{\lambda}_1 < 1$. Hence, if $\nu^2 \leq \sigma_1^2$, then we also have $\sigma_2^2 = \hat{\lambda}_1 \nu^2 < \sigma_1^2$ and $\lambda_2 = \frac{\sigma_2^2}{\sigma_2^2 + \nu^2} < \frac{\sigma_1^2}{\sigma_1^2 + \nu^2} = \lambda_1$. Recall from the proof of Part I that $\hat{\lambda}_t$ monotonically increases in λ_t . We thus conclude that $\hat{\lambda}_2 < \hat{\lambda}_1$, and the claim follows from Part I.

Part III. First, we show that *if $\alpha = 0$, then the sequence $(\hat{\lambda}_t)_{t \in \mathbb{N}}$ is decreasing*. We observe that

$$\frac{\partial f}{\partial \hat{\lambda}_t} \Big|_{\alpha=0} = -\frac{\gamma}{2} \nu^2 \sum_{\tau=t+1}^{\infty} \phi_{\tau-t} \frac{1}{(1+(\tau-(t+1))\hat{\lambda}_t)^2} + \nu^2 \sum_{\tau=t+1}^{\infty} \delta_{\tau-t} \frac{(\tau-t)}{(1+(\tau-(t+1))\hat{\lambda}_t)^3} \frac{\lambda_t - \hat{\lambda}_t}{1 - \lambda_t}.$$

is strictly negative for any $\hat{\lambda}_t \geq \lambda_t$. Hence, if $\alpha = 0$, then $\hat{\lambda}_t < \lambda_t$ for any $t \geq 1$. In particular, $\sigma_2^2 = \hat{\lambda}_1 \nu^2 < \lambda_1 \nu^2 = \sigma_1^2(1 - \lambda_1) < \sigma_1^2$ and thus $\lambda_2 < \lambda_1$. Because $\hat{\lambda}_t$ monotonically increases in λ_t , we thus have $\hat{\lambda}_2 < \hat{\lambda}_1$. The claim thus follows from Part I.

Second, fixing $\lambda_t \in [0, 1)$, we show that $\hat{\lambda}_t$ monotonically increases in α . For any $\hat{\lambda}_t < 1$,

$$\frac{\partial f}{\partial \alpha \partial \hat{\lambda}_t} = 2\alpha \gamma^2 \Phi^2 \frac{\sum_{\tau=t+1}^{\infty} \delta_{\tau-t}(\tau-t) \frac{(1-\hat{\lambda}_t)}{(1+(\tau-(t+1))\hat{\lambda}_t)^3}}{\left(1 + \sum_{\tau=t+1}^{\infty} \delta_{\tau-t} \left(\frac{1-\hat{\lambda}_t}{1+(\tau-(t+1))\hat{\lambda}_t}\right)^2\right)^2} > 0.$$

As argued above, the claim then immediately follows from Theorem 10.6 in Sundaram (1996).

Third, we show that *there exists some $\underline{\alpha} \in \mathbb{R}$, such that for any $\alpha > \underline{\alpha}$, $(\hat{\lambda}_t)_{t \in \mathbb{N}}$ is increasing.* Notice that, for any fixed $\hat{\lambda}_t < 1$, we have $\lim_{\alpha \rightarrow \infty} \frac{\partial f}{\partial \hat{\lambda}_t} = \infty$. This implies $\lim_{\alpha \rightarrow \infty} \hat{\lambda}_t = 1$. Now fix an initial variance $\sigma_1^2 < \nu^2$, and recall that $\hat{\lambda}_1$ monotonically increases in α . Hence, there exists some $\underline{\alpha} \in \mathbb{R}$, such that for any $\alpha > \underline{\alpha}$, we have $\sigma_2^2 = \hat{\lambda}_1 \nu^2 > \sigma_1^2$, which in turn implies $\lambda_2 > \lambda_1$. Because $\hat{\lambda}_t$ monotonically increases in λ_t , we thus have $\hat{\lambda}_2 > \hat{\lambda}_1$. The claim follows from Part I. \square

Proof of Proposition 3. Part I. Fix $\gamma, \Phi, \Delta, \Omega \in \mathbb{R}_{>0}$. For the sake of a contradiction, suppose that $\hat{\lambda}_t > 0$ for any $t \in \mathbb{N}$ with $\lim_{t \rightarrow \infty} \hat{\lambda}_t = 0$. By Lemma 3, $(\hat{\lambda}_t)_{t \in \mathbb{N}}$ must be decreasing.

First, suppose that $\alpha^2 < \frac{(1+\Delta)^2 \nu^2}{\Omega 2\gamma\Phi}$. By our assumption towards a contradiction, we can find some $t' \in \mathbb{N}$ such that for any $t \geq t'$,

$$\alpha^2 < \frac{(1+\Delta)^2 \nu^2}{\Omega 2\gamma\Phi} - \hat{\lambda}_{t-1} \frac{\nu^2}{\gamma^2 \Phi^2} (1+\Delta)^2.$$

For any $t \geq t'$, we have

$$\begin{aligned} \left. \frac{\partial f}{\partial \hat{\lambda}_t} \right|_{\hat{\lambda}_t=0} &= \frac{\Omega}{(1+\Delta)^2} \alpha^2 \gamma^2 \Phi^2 - \frac{\gamma}{2} \nu^2 \Phi + \nu^2 \frac{\lambda_t}{1-\lambda_t} \Omega \\ &= \frac{\Omega}{(1+\Delta)^2} \alpha^2 \gamma^2 \Phi^2 - \frac{\gamma}{2} \nu^2 \Phi + \nu^2 \hat{\lambda}_{t-1} \Omega \\ &< 0. \end{aligned}$$

Moreover, because $\frac{\partial f}{\partial \hat{\lambda}_t}$ is continuous in $\hat{\lambda}_t$, we can find some $\epsilon_t > 0$, such that $\frac{\partial f}{\partial \hat{\lambda}_t} < 0$ for any $\hat{\lambda}_t \in [0, \epsilon_t)$. Hence, $\hat{\lambda}_t = 0$ or $\hat{\lambda}_t \geq \epsilon_t$. Because (i) $\frac{\partial^2 f}{\partial \lambda_t \partial \hat{\lambda}_t} > 0$ for any $\hat{\lambda}_t < 1$, (ii) $\lambda_{t+1} = \frac{\hat{\lambda}_t}{\lambda_{t+1}}$ is increasing in $\hat{\lambda}_t$, and (iii) $(\hat{\lambda}_t)_{t \in \mathbb{N}}$ is decreasing, the sequence $(\epsilon_t)_{t \geq t'}$ must be increasing. By our

assumption towards a contradiction, we can thus find a period $t'' \geq t'$ such that $\hat{\lambda}_{t''-1} < \epsilon_{t'} \leq \epsilon_{t''}$. Hence, for any $t \geq t''$, the agent chooses $\hat{\lambda}_t = 0$; a contradiction.

Second, suppose that $\alpha^2 > \frac{(1+\Delta)^2}{\Omega} \frac{\nu^2}{2\gamma\Phi}$. Then, for any t , we have

$$\left. \frac{\partial f}{\partial \hat{\lambda}_t} \right|_{\hat{\lambda}_t=0} = \frac{\Omega}{(1+\Delta)^2} \alpha^2 \gamma^2 \Phi^2 - \frac{\gamma}{2} \nu^2 \Phi + \nu^2 \frac{\lambda_t}{1-\lambda_t} \Omega > 0.$$

Because $\frac{\partial f}{\partial \hat{\lambda}_t}$ is continuous in $\hat{\lambda}_t$, we can find some $\epsilon_t > 0$, such that $\frac{\partial f}{\partial \hat{\lambda}_t} > 0$ for any $\hat{\lambda}_t \in [0, \epsilon_t)$. Hence, it has to be true that $\hat{\lambda}_t \geq \epsilon_t$. Because (i) $\frac{\partial^2 f}{\partial \lambda_t \partial \hat{\lambda}_t} > 0$ for any $\hat{\lambda}_t < 1$, (ii) $\lambda_{t+1} = \frac{\hat{\lambda}_t}{\hat{\lambda}_t + 1}$ is increasing in $\hat{\lambda}_t$, and (iii) $(\hat{\lambda}_t)_{t \in \mathbb{N}}$ is decreasing, the sequence $(\epsilon_t)_{t \in \mathbb{N}}$ must be decreasing. Still, since $\alpha^2 > \frac{(1+\Delta)^2}{\Omega} \frac{\nu^2}{2\gamma\Phi}$, $\lim_{t \rightarrow \infty} \epsilon_t = \epsilon > 0$. This implies that $\lim_{t \rightarrow \infty} \hat{\lambda}_t \geq \epsilon > 0$; a contradiction.

Part II. We first show that *if $\alpha = 0$, the agent develops a dogma*. As we have argued in the proof of Part III of Lemma 3, if $\alpha = 0$, then $\hat{\lambda}_t < \lambda_t$ for any $t \geq 1$. We bound $\hat{\lambda}_t$ from above by the weight λ_t^b that a fully Bayesian agent starting out with the same prior belief, $\mathcal{N}(\mu_1, \sigma_1^2)$, would put on signal s_t . Formally, we show that, *for any t , we have $\hat{\lambda}_t < \lambda_t^b$* .

We prove this claim by induction (over t). In period 1, we have $\hat{\lambda}_1 < \lambda_1 = \lambda_1^b$. Now assume that $\hat{\lambda}_t < \lambda_t^b$ holds. We have to show that also $\hat{\lambda}_{t+1} < \lambda_{t+1}^b$. Since $\hat{\lambda}_t < \lambda_t$ for any $t \geq 1$, we obtain

$$\hat{\lambda}_{t+1} < \lambda_{t+1} = \frac{\hat{\lambda}_t}{\hat{\lambda}_t + 1} < \frac{\lambda_t^b}{\lambda_t^b + 1} = \lambda_{t+1}^b,$$

where the second inequality follows from the induction hypothesis.

Because $\hat{\lambda}_t < \lambda_t^b$, we further conclude that

$$\hat{\lambda}_t < \frac{\lambda_1}{1 + (t-1)\lambda_1} \rightarrow 0 \quad \text{as } t \rightarrow \infty.$$

By Part I, therefore, there exists some $t' \in \mathbb{N}$, such that for any $t \geq t'$, we have $\hat{\lambda}_t = 0$.

Second, by the proof of Part I, an agent with $\alpha^2 > \frac{(1+\Delta)^2}{\Omega} \frac{\nu^2}{2\gamma\Phi}$ does not develop a dogma.

Third, we show that, for a given λ_1 , the weight $\hat{\lambda}_t$ monotonically increases in α . As we have argued in the proof of Lemma 3, fixing λ_t , the weight $\hat{\lambda}_t$ indeed monotonically increases in α . This implies that $\hat{\lambda}_1$ monotonically increases in α . Next, we prove our claim by induction (over t). Compare α' and $\alpha'' > \alpha'$. By the induction hypothesis, we have $\hat{\lambda}_t(\alpha'') > \hat{\lambda}_t(\alpha')$, which in turn implies that $\lambda'' := \lambda_{t+1}(\alpha'') > \lambda_{t+1}(\alpha') =: \lambda'$. Again, as we have argued in the proof of Lemma 3, fixing α , the weight $\hat{\lambda}_t$ increases in λ_t . Hence, we have $\hat{\lambda}_{t+1}(\alpha'', \lambda'') \geq \hat{\lambda}_{t+1}(\alpha'', \lambda') \geq \hat{\lambda}_{t+1}(\alpha', \lambda')$.

Fourth, we argue that this monotonicity is preserved in the limit. Consider α' and $\alpha'' > \alpha'$. Let $(\hat{\lambda}'_t)_{t \in \mathbb{N}}$ and $(\hat{\lambda}''_t)_{t \in \mathbb{N}}$ be the corresponding sequences, with limit weights $\hat{\lambda}'_\infty$ and $\hat{\lambda}''_\infty$, respectively. Define $d_t = \hat{\lambda}''_t - \hat{\lambda}'_t \geq 0$, with a limit $d_\infty = \hat{\lambda}''_\infty - \hat{\lambda}'_\infty$. For the sake of a contradiction, suppose $d_\infty < 0$, and set $\epsilon = -d_\infty/2$. Because $(d_t)_{t \in \mathbb{N}}$ converges, there exists some t^* , so that for any $t \geq t^*$,

$$|d_t - d_\infty| < \epsilon \quad \text{or, equivalently,} \quad d_\infty - \epsilon < d_t < d_\infty + \epsilon = \frac{d_\infty}{2} < 0;$$

a contradiction. Hence, $d_\infty \geq 0$ and thus $\hat{\lambda}''_\infty \geq \hat{\lambda}'_\infty$. In sum, we conclude that, there exists some $\bar{\alpha} \in \mathbb{R}_{>0}$, such that the agent develops a dogma if and only if $\alpha \leq \bar{\alpha}$. \square

Proof of Proposition 4. Parts I and II. By Proposition 3, if $\alpha < \bar{\alpha}$, the agent develops a dogma. Once the agent has developed a dogma, she does no longer react to the signals. Still, since she treats μ_t as truth, in every period t large enough, she increases her mean belief by $\frac{\alpha\gamma\Phi}{1+\Delta}$. In particular, if $\alpha = 0$, the agent does not bias her mean belief. Part I then follows directly from Lemma 2 with

$$\omega_E := \sum_{t=1}^{\tau^*-1} \hat{\lambda}_t \prod_{\ell=t+1}^{\tau^*-1} (1 - \hat{\lambda}_\ell) \quad \text{and} \quad \omega_V := \sum_{t=1}^{\tau^*-1} \hat{\lambda}_t^2 \prod_{\ell=t+1}^{\tau^*-1} (1 - \hat{\lambda}_\ell)^2,$$

where $\tau^* \in \mathbb{N}$ is the period in which the agent chooses to be dogmatic.

Part III. Let $\alpha > \bar{\alpha}$. By Lemma 2, the agent's mean belief in period τ is given by

$$\hat{\mu}_\tau = \mu_1 \prod_{t=1}^{\tau-1} (1 - \hat{\lambda}_t) + \sum_{t=1}^{\tau} \hat{b}_t \prod_{\ell=t}^{\tau-1} (1 - \hat{\lambda}_\ell) + \sum_{t=1}^{\tau-1} s_t \hat{\lambda}_t \prod_{\ell=t+1}^{\tau-1} (1 - \hat{\lambda}_\ell). \quad (9)$$

In the following, we proceed in three steps. First, we characterize the agent's expected long-run mean belief. In a second step, we derive the variance of the agent's long-run mean belief. Finally, we argue that the agent's long-run mean belief, μ_∞ , is normally distributed.

1. Step. We first determine the agent's expected long-run mean belief. When taking the expectation over the sequence of signals in Eq. (9), we obtain

$$\mathbb{E}[\hat{\mu}_\tau] = \mu_1 \prod_{t=1}^{\tau-1} (1 - \hat{\lambda}_t) + \sum_{t=1}^{\tau} \hat{b}_t \prod_{\ell=t}^{\tau-1} (1 - \hat{\lambda}_\ell) + \theta \sum_{t=1}^{\tau-1} \hat{\lambda}_t \prod_{\ell=t+1}^{\tau-1} (1 - \hat{\lambda}_\ell). \quad (10)$$

We study the limit behavior of the three terms on the right-hand side in Eq. (10) separately.

Lemma 4 (Limit Weight on the Prior).

$$\lim_{\tau \rightarrow \infty} \mu_1 \prod_{t=1}^{\tau-1} (1 - \hat{\lambda}_t) = 0.$$

Proof. Let $\underline{\lambda} := \inf_{t \in \mathbb{N}} \hat{\lambda}_t$, and recall that, by the preliminaries, $\underline{\lambda} > 0$. This implies that

$$|\mu_1| \prod_{t=1}^{\tau-1} (1 - \hat{\lambda}_t) < |\mu_1| (1 - \underline{\lambda})^{\tau-1} \longrightarrow 0 \quad \text{as } \tau \rightarrow \infty,$$

which in turn yields the claim. \square

Lemma 5 (Limit Bias).

$$\lim_{\tau \rightarrow \infty} \sum_{t=1}^{\tau} \hat{b}_t \prod_{\ell=t}^{\tau-1} (1 - \hat{\lambda}_\ell) = \frac{\hat{b}_\infty}{\hat{\lambda}_\infty}.$$

Proof. For the sake of a contradiction, suppose that there exists some $\epsilon > 0$ such that

$$\left| \frac{\hat{b}_\infty}{\hat{\lambda}_\infty} - \lim_{\tau \rightarrow \infty} \sum_{t=1}^{\tau} \hat{b}_t \prod_{\ell=t}^{\tau-1} (1 - \hat{\lambda}_\ell) \right| > \epsilon.$$

Because $\lim_{t \rightarrow \infty} \hat{\lambda}_t = \hat{\lambda}_\infty$ and $\lim_{t \rightarrow \infty} \hat{b}_t = \hat{b}_\infty$ exist, we can find some τ' , such that for all $\tau > \tau'$,

$$\left| \sum_{t=\tau'+1}^{\tau} \hat{b}_t \prod_{\ell=t}^{\tau-1} (1 - \hat{\lambda}_\ell) - \sum_{t=\tau'+1}^{\tau} \hat{b}_\infty \prod_{\ell=t}^{\tau-1} (1 - \hat{\lambda}_\infty) \right| < \epsilon. \quad (11)$$

Now fix such an τ' and choose $\tau > \tau' + 1$. Then,

$$\begin{aligned} \sum_{t=1}^{\tau} \hat{b}_t \prod_{\ell=t}^{\tau-1} (1 - \hat{\lambda}_\ell) &= \sum_{t=1}^{\tau'} \hat{b}_t \prod_{\ell=t}^{\tau-1} (1 - \hat{\lambda}_\ell) + \sum_{t=\tau'+1}^{\tau} \hat{b}_\infty \prod_{\ell=t}^{\tau-1} (1 - \hat{\lambda}_\infty) \\ &\quad + \sum_{t=\tau'+1}^{\tau} \hat{b}_t \prod_{\ell=t}^{\tau-1} (1 - \hat{\lambda}_\ell) - \sum_{t=\tau'+1}^{\tau} \hat{b}_\infty \prod_{\ell=t}^{\tau-1} (1 - \hat{\lambda}_\infty). \end{aligned} \quad (12)$$

Again using the fact that $\underline{\lambda} > 0$, we observe that

$$\begin{aligned} \sum_{t=1}^{\tau'} \hat{b}_t \prod_{\ell=t}^{\tau-1} (1 - \hat{\lambda}_\ell) &= \sum_{t=1}^{\tau'} \hat{b}_t \prod_{\ell=t}^{\tau'-1} (1 - \hat{\lambda}_\ell) \prod_{\ell=\tau'}^{\tau-1} (1 - \hat{\lambda}_\ell) \\ &< (1 - \underline{\lambda})^{\tau-\tau'} \sum_{t=1}^{\tau'} \hat{b}_t \prod_{\ell=t}^{\tau'-1} (1 - \hat{\lambda}_\ell) \longrightarrow 0 \quad \text{as } \tau \rightarrow \infty. \end{aligned}$$

Moreover, we have

$$\begin{aligned} \sum_{t=\tau'+1}^{\tau} \hat{b}_{\infty} \prod_{\ell=t}^{\tau-1} (1 - \hat{\lambda}_{\infty}) &= \hat{b}_{\infty} \sum_{t=\tau'+1}^{\tau} (1 - \hat{\lambda}_{\infty})^{\tau-t} \\ &= \hat{b}_{\infty} \sum_{\ell=0}^{\tau-(\tau'+1)} (1 - \hat{\lambda}_{\infty})^{\ell} \longrightarrow \frac{\hat{b}_{\infty}}{\hat{\lambda}_{\infty}} \quad \text{as } \tau \rightarrow \infty. \end{aligned}$$

Combining these two observations with Eq. (12), we conclude that

$$\lim_{\tau \rightarrow \infty} \sum_{t=1}^{\tau} \hat{b}_t \prod_{\ell=t}^{\tau-1} (1 - \hat{\lambda}_{\ell}) = \frac{\hat{b}_{\infty}}{\hat{\lambda}_{\infty}} + \lim_{\tau \rightarrow \infty} \left\{ \sum_{t=\tau'+1}^{\tau} \hat{b}_t \prod_{\ell=t}^{\tau-1} (1 - \hat{\lambda}_{\ell}) - \sum_{t=\tau'+1}^{\tau} \hat{b}_{\infty} \prod_{\ell=t}^{\tau-1} (1 - \hat{\lambda}_{\infty}) \right\}.$$

Using (11), it therefore follows

$$\left| \frac{\hat{b}_{\infty}}{\hat{\lambda}_{\infty}} - \lim_{\tau \rightarrow \infty} \sum_{t=1}^{\tau} \hat{b}_t \prod_{\ell=t}^{\tau-1} (1 - \hat{\lambda}_{\ell}) \right| < \epsilon;$$

a contradiction. □

Lemma 6 (Limit Weight on the Truth).

$$\sum_{t=1}^{\tau-1} \hat{\lambda}_t \prod_{\ell=t+1}^{\tau-1} (1 - \hat{\lambda}_{\ell}) = 1.$$

Proof. For the sake of a contradiction, suppose that there exists some $\epsilon > 0$ such that

$$\left| 1 - \lim_{\tau \rightarrow \infty} \sum_{t=1}^{\tau-1} \hat{\lambda}_t \prod_{\ell=t+1}^{\tau-1} (1 - \hat{\lambda}_{\ell}) \right| > \epsilon.$$

Because $\lim_{t \rightarrow \infty} \hat{\lambda}_t = \hat{\lambda}_{\infty}$ and $\lim_{t \rightarrow \infty} \hat{b}_t = \hat{b}_{\infty}$ exist, we can find some τ' , such that for all $\tau > \tau'$,

$$\left| \sum_{t=\tau'+1}^{\tau} \hat{\lambda}_t \prod_{\ell=t+1}^{\tau-1} (1 - \hat{\lambda}_{\ell}) - \sum_{t=\tau'+1}^{\tau} \hat{\lambda}_{\infty} \prod_{\ell=t+1}^{\tau-1} (1 - \hat{\lambda}_{\infty}) \right| < \epsilon. \quad (13)$$

Now fix such a τ' and choose $\tau > \tau' + 2$. Then,

$$\begin{aligned} \sum_{t=1}^{\tau-1} \hat{\lambda}_t \prod_{\ell=t+1}^{\tau-1} (1 - \hat{\lambda}_{\ell}) &= \sum_{t=1}^{\tau'} \hat{\lambda}_t \prod_{\ell=t+1}^{\tau-1} (1 - \hat{\lambda}_{\ell}) + \sum_{t=\tau'+1}^{\tau-1} \hat{\lambda}_{\infty} \prod_{\ell=t+1}^{\tau-1} (1 - \hat{\lambda}_{\infty}) \\ &\quad + \sum_{t=\tau'+1}^{\tau-1} \hat{\lambda}_t \prod_{\ell=t+1}^{\tau-1} (1 - \hat{\lambda}_{\ell}) - \sum_{t=\tau'+1}^{\tau-1} \hat{\lambda}_{\infty} \prod_{\ell=t+1}^{\tau-1} (1 - \hat{\lambda}_{\infty}). \end{aligned} \quad (14)$$

Again using the fact that $\underline{\lambda} > 0$, we observe that

$$\begin{aligned} \sum_{t=1}^{\tau'} \hat{\lambda}_t \prod_{\ell=t+1}^{\tau-1} (1 - \hat{\lambda}_\ell) &= \sum_{t=1}^{\tau'} \hat{\lambda}_t \prod_{\ell=t+1}^{\tau'-1} (1 - \hat{\lambda}_\ell) \prod_{\ell=\tau'}^{\tau-1} (1 - \hat{\lambda}_\ell) \\ &< (1 - \underline{\lambda})^{\tau-\tau'} \sum_{t=1}^{\tau'} \hat{\lambda}_t \prod_{\ell=t+1}^{\tau'-1} (1 - \hat{\lambda}_\ell) \longrightarrow 0 \quad \text{as } \tau \rightarrow \infty. \end{aligned}$$

Moreover, we have

$$\begin{aligned} \sum_{t=\tau'+1}^{\tau-1} \hat{\lambda}_\infty \prod_{\ell=t+1}^{\tau-1} (1 - \hat{\lambda}_\infty) &= \hat{\lambda}_\infty \sum_{t=\tau'+1}^{\tau-1} (1 - \hat{\lambda}_\infty)^{\tau-(t+1)} \\ &= \hat{\lambda}_\infty \sum_{\ell=0}^{\tau-(\tau'+2)} (1 - \hat{\lambda}_\infty)^\ell \longrightarrow 1 \quad \text{as } \tau \rightarrow \infty. \end{aligned}$$

Combining these two observations with Eq. (14), we conclude that

$$\lim_{\tau \rightarrow \infty} \sum_{t=1}^{\tau-1} \hat{\lambda}_t \prod_{\ell=t+1}^{\tau-1} (1 - \hat{\lambda}_\ell) = 1 + \lim_{\tau \rightarrow \infty} \left\{ \sum_{t=\tau'+1}^{\tau-1} \hat{\lambda}_t \prod_{\ell=t+1}^{\tau-1} (1 - \hat{\lambda}_\ell) - \sum_{t=\tau'+1}^{\tau-1} \hat{\lambda}_\infty \prod_{\ell=t+1}^{\tau-1} (1 - \hat{\lambda}_\infty) \right\}.$$

Using (13), it therefore follows

$$\left| 1 - \lim_{\tau \rightarrow \infty} \sum_{t=1}^{\tau-1} \hat{\lambda}_t \prod_{\ell=t+1}^{\tau-1} (1 - \hat{\lambda}_\ell) \right| < \epsilon;$$

a contradiction. □

Combining Lemma 4, 5, and 6, we obtain

$$\lim_{\tau \rightarrow \infty} \mathbb{E}[\mu_\tau] = \theta + \frac{\hat{b}_\infty}{\hat{\lambda}_\infty}.$$

2. *Step.* We next derive the variance of the agent's long-run mean belief. By Eq. (9), we have

$$\text{Var}(\hat{\mu}_\tau) = \nu^2 \sum_{t=1}^{\tau-1} \hat{\lambda}_t^2 \prod_{\ell=t+1}^{\tau-1} (1 - \hat{\lambda}_\ell)^2.$$

By the exact same arguments as in the proof of Lemma 6, we have

$$\lim_{\tau \rightarrow \infty} \sum_{t=1}^{\tau-1} \hat{\lambda}_t^2 \prod_{\ell=t+1}^{\tau-1} (1 - \hat{\lambda}_\ell)^2 = \hat{\lambda}_\infty^2 \frac{1}{1 - (1 - \hat{\lambda}_\infty)^2} = \frac{\hat{\lambda}_\infty}{2 - \hat{\lambda}_\infty}.$$

Hence, we have

$$\lim_{\tau \rightarrow \infty} \text{Var}(\hat{\mu}_\tau) = \nu^2 \frac{\hat{\lambda}_\infty}{2 - \hat{\lambda}_\infty}.$$

3. *Step.* We finally argue that the agent's long-run mean belief is normally distributed. Taking the limit of Eq. (9), we have

$$\hat{\mu}_\infty - \frac{\hat{b}_\infty}{\hat{\lambda}_\infty} = \sum_{t=1}^{\infty} s_t \hat{\lambda}_t \underbrace{\prod_{\ell=t+1}^{\infty} (1 - \hat{\lambda}_\ell)}_{=: \tilde{s}_t},$$

which is an infinite sum of independent and normally distributed random variables

$$\tilde{s}_t \sim \mathcal{N} \left(\theta \hat{\lambda}_t \prod_{\ell=t+1}^{\infty} (1 - \hat{\lambda}_\ell), \nu^2 \hat{\lambda}_t^2 \prod_{\ell=t+1}^{\infty} (1 - \hat{\lambda}_\ell)^2 \right).$$

As we have seen in the first two steps above, both

$$\sum_{t=1}^{\infty} \mathbb{E}[\tilde{s}_t] = \theta \quad \text{and} \quad \sum_{t=1}^{\infty} \text{Var}(\tilde{s}_t) = \nu^2 \frac{\hat{\lambda}_\infty}{2 - \hat{\lambda}_\infty}$$

exist. Hence, because the random variables $\{\tilde{s}_t\}_{t \in \mathbb{N}}$ are independently and normally distributed, the characteristic function¹⁴ of $\sum_{t=1}^{\tau} \tilde{s}_t$ converges pointwise to

$$\exp \left(-\frac{t^2}{2} \sum_{t=1}^{\infty} \text{Var}(\tilde{s}_t) + it \sum_{t=1}^{\infty} \mathbb{E}[\tilde{s}_t] \right), \quad \text{as } \tau \rightarrow \infty.$$

By Lévy's convergence theorem, therefore, $\sum_{t=1}^{\tau} \tilde{s}_t$ converges in distribution to

$$\mathcal{N} \left(\theta, \nu^2 \frac{\hat{\lambda}_\infty}{2 - \hat{\lambda}_\infty} \right), \quad \text{as } \tau \rightarrow \infty,$$

and, as a consequence, μ_∞ converges in distribution to

$$\mathcal{N} \left(\theta + \frac{\hat{b}_\infty}{\hat{\lambda}_\infty}, \nu^2 \frac{\hat{\lambda}_\infty}{2 - \hat{\lambda}_\infty} \right), \quad \text{as } \tau \rightarrow \infty.$$

This completes the proof. □

¹⁴See, for example, [https://en.wikipedia.org/wiki/Characteristic_function_\(probability_theory\)](https://en.wikipedia.org/wiki/Characteristic_function_(probability_theory)) for a definition and properties of the characteristic function (accessed on February 29, 2024).

Proof of Corollary 2. Part I. As we have argued in the proof of Proposition 4, $\hat{\lambda}_\infty$ monotonically increases in α . Next, for the sake of a contradiction, suppose $\lim_{\alpha \rightarrow \infty} \hat{\lambda}_\infty < 1$. Because $\hat{\lambda}_\infty > 0$,

$$\begin{aligned} 0 &= \left. \frac{\partial f}{\partial \hat{\lambda}_t} \right|_{\hat{\lambda}_t = \hat{\lambda}_\infty, \lambda_t = \hat{\lambda}_\infty / (1 + \hat{\lambda}_\infty)} \\ &= \alpha^2 \gamma^2 \Phi^2 \frac{\sum_{\tau=t+1}^{\infty} \delta_{\tau-t} (\tau-t) \frac{(1-\hat{\lambda}_\infty)}{(1+(\tau-(t+1))\hat{\lambda}_\infty)^3}}{\left(1 + \sum_{\tau=t+1}^{\infty} \delta_{\tau-t} \left(\frac{1-\hat{\lambda}_\infty}{1+(\tau-(t+1))\hat{\lambda}_\infty}\right)^2\right)^2} - \frac{\gamma}{2} \nu^2 \sum_{\tau=t+1}^{\infty} \phi_{\tau-t} \frac{1}{(1+(\tau-(t+1))\hat{\lambda}_\infty)^2} \\ &\quad - \hat{\lambda}_\infty^2 \nu^2 \sum_{\tau=t+1}^{\infty} \delta_{\tau-t} \frac{(\tau-t)}{(1+(\tau-(t+1))\hat{\lambda}_\infty)^3}, \end{aligned}$$

with f being defined as in the proof of Lemma 1. Because the second term is bounded from below by $-\frac{\gamma}{2} \nu^2 \Phi > -\infty$ and the third term is bounded from below by $-\nu^2 \Omega > -\infty$, we have

$$\frac{\sum_{\tau=t+1}^{\infty} \delta_{\tau-t} (\tau-t) \frac{(1-\hat{\lambda}_\infty)}{(1+(\tau-(t+1))\hat{\lambda}_\infty)^3}}{\left(1 + \sum_{\tau=t+1}^{\infty} \delta_{\tau-t} \left(\frac{1-\hat{\lambda}_\infty}{1+(\tau-(t+1))\hat{\lambda}_\infty}\right)^2\right)^2} \rightarrow 0 \quad \text{as } \alpha \rightarrow \infty,$$

for otherwise $\left. \frac{\partial f}{\partial \hat{\lambda}_t} \right|_{\hat{\lambda}_t = \hat{\lambda}_\infty, \lambda_t = \hat{\lambda}_\infty / (1 + \hat{\lambda}_\infty)} \rightarrow \infty$ as $\alpha \rightarrow \infty$. Hence, $\hat{\lambda}_\infty \rightarrow 1$ as $\alpha \rightarrow \infty$; a contradiction.

Part II. For any $\alpha < \bar{\alpha}$, by Proposition 4, $\mathbb{E}[\hat{\mu}_\infty] = \infty$. For any $\alpha > \bar{\alpha}$, by Proposition 4,

$$\mathbb{E}[\hat{\mu}_\infty] = \theta + \frac{b(\hat{\lambda}_\infty)}{\hat{\lambda}_\infty} \geq \theta + b(\hat{\lambda}_\infty) \geq \theta + \frac{\alpha \gamma \Phi}{1 + \Delta},$$

where the first inequality follows from $\hat{\lambda}_\infty \leq 1$, and the second inequality follows from Lemma 1 and the fact that $b(\cdot)$ is increasing. It follows that $\lim_{\alpha \rightarrow \infty} \mathbb{E}[\hat{\mu}_\infty] = \infty$.

Part III. Let $F(\cdot)$ be the CDF of the standard-normal distribution, and pick some $\hat{\mu} \in \mathbb{R}$. Because $\hat{\lambda}_\infty > 0$ for any $\alpha > \bar{\alpha}$, and because $s_t \sim \mathcal{N}(\theta, \nu^2)$, we have

$$\mathbb{P}[(1 - \hat{\lambda}_\infty)\hat{\mu} + \hat{\lambda}_\infty s_t < \theta] = \mathbb{P}\left[s_t < \frac{\theta - (1 - \hat{\lambda}_\infty)\hat{\mu}}{\hat{\lambda}_\infty}\right] = F\left(-\frac{(1 - \hat{\lambda}_\infty)(\theta - \hat{\mu})}{\hat{\lambda}_\infty \nu}\right).$$

The claim follows from the fact that $F(x) > 0$ for all $x \in \mathbb{R}$. □

Proof of Corollary 3. Fix any $\sigma_1^2 \in \mathbb{R}_{\geq 0}$, $\alpha \in \mathbb{R}_{> 0}$, and $\gamma \in \mathbb{R}_{\geq 0}$.

Part I. We first show that $\hat{\lambda}_1 = 0$ for any $\nu^2 > \frac{2\Omega}{\gamma\phi_1}(\alpha^2\gamma^2\Phi^2 + \sigma_1^2)$. Notice that, for any $\hat{\lambda}_1 \in [0, 1]$,

$$\begin{aligned} \frac{\partial f}{\partial \hat{\lambda}_1} &= \alpha^2\gamma^2\Phi^2 \frac{\sum_{\tau=2}^{\infty} \delta_{\tau-1}(\tau-1) \frac{(1-\hat{\lambda}_1)}{(1+(\tau-(t+1))\hat{\lambda}_1)^3}}{\left(1 + \sum_{\tau=2}^{\infty} \delta_{\tau-1} \left(\frac{1-\hat{\lambda}_1}{1+(\tau-2)\hat{\lambda}_1}\right)^2\right)^2} - \frac{\gamma}{2}\nu^2 \sum_{\tau=2}^{\infty} \phi_{\tau-1} \frac{1}{(1+(\tau-2)\hat{\lambda}_1)^2} \\ &\quad + \nu^2 \sum_{\tau=2}^{\infty} \delta_{\tau-1} \frac{(\tau-1)}{(1+(\tau-2)\hat{\lambda}_1)^3} \frac{\lambda_1 - \hat{\lambda}_1}{1-\lambda_1} \\ &< \alpha^2\gamma^2\Phi^2\Omega - \frac{\gamma}{2}\nu^2\phi_1 + \sigma_1^2\Omega, \end{aligned}$$

which is strictly negative for any $\nu^2 > \frac{2\Omega}{\gamma\phi_1}(\alpha^2\gamma^2\Phi^2 + \sigma_1^2)$. As a result, for any $\nu^2 > \frac{2\Omega}{\gamma\phi_1}(\alpha^2\gamma^2\Phi^2 + \sigma_1^2)$, the agent to chooses $\hat{\lambda}_1 = 0$. By Lemma 3, and by the fact that $\hat{\lambda}_t \geq 0$, it follows that $\hat{\lambda}_{\infty} = 0$. With this, by Proposition 4, we also have $\mathbb{E}[\hat{\mu}_{\infty}] = \infty$ for any $\nu^2 > \frac{2\Omega}{\gamma\phi_1}(\alpha^2\gamma^2\Phi^2 + \sigma_1^2)$.

Part II. We first show that $\lim_{\nu^2 \rightarrow 0} \hat{\lambda}_{\infty} = 1$ for any $\alpha > 0$. For the sake of a contradiction, suppose that $\lim_{\nu^2 \rightarrow 0} \hat{\lambda}_{\infty} < 1$, and notice that $\hat{\lambda}_{\infty}$ must satisfy

$$\begin{aligned} 0 &\geq \frac{\partial f}{\partial \hat{\lambda}_t} \Big|_{\hat{\lambda}_t = \hat{\lambda}_{\infty}, \lambda_t = \hat{\lambda}_{\infty}/(1+\hat{\lambda}_{\infty})} \\ &= \alpha^2\gamma^2\Phi^2 \frac{\sum_{\tau=t+1}^{\infty} \delta_{\tau-t}(\tau-t) \frac{(1-\hat{\lambda}_{\infty})}{(1+(\tau-(t+1))\hat{\lambda}_{\infty})^3}}{\left(1 + \sum_{\tau=t+1}^{\infty} \delta_{\tau-t} \left(\frac{1-\hat{\lambda}_{\infty}}{1+(\tau-(t+1))\hat{\lambda}_{\infty}}\right)^2\right)^2} - \frac{\gamma}{2}\nu^2 \sum_{\tau=t+1}^{\infty} \phi_{\tau-t} \frac{1}{(1+(\tau-(t+1))\hat{\lambda}_{\infty})^2} \\ &\quad - \hat{\lambda}_{\infty}^2 \nu^2 \sum_{\tau=t+1}^{\infty} \delta_{\tau-t} \frac{(\tau-t)}{(1+(\tau-(t+1))\hat{\lambda}_{\infty})^3}. \end{aligned}$$

As $\nu^2 \rightarrow 0$, the right-hand side above approaches

$$\lim_{\nu^2 \rightarrow 0} \left\{ \alpha^2\gamma^2\Phi^2 \frac{\sum_{\tau=t+1}^{\infty} \delta_{\tau-t}(\tau-t) \frac{(1-\hat{\lambda}_{\infty})}{(1+(\tau-(t+1))\hat{\lambda}_{\infty})^3}}{\left(1 + \sum_{\tau=t+1}^{\infty} \delta_{\tau-t} \left(\frac{1-\hat{\lambda}_{\infty}}{1+(\tau-(t+1))\hat{\lambda}_{\infty}}\right)^2\right)^2} \right\},$$

which is strictly positive by our assumption towards a contradiction. Hence, the agent prefers to increase $\hat{\lambda}_{\infty}$; a contradiction. With this, by Proposition 4, we further have $\mathbb{E}[\hat{\mu}_{\infty}] = \theta + \alpha\gamma\Phi$. \square

Proof of Corollary 4. Notice that f , as defined in the proof of Lemma 1, does not depend on μ_t . Hence, also $\hat{\lambda}_t$ is independent of μ_t , and as a result, $\hat{\lambda}_{\infty}$ is independent of μ_1 . Moreover, as we

have shown in the proof of Lemma 3, $\hat{\lambda}_t$ (weakly) increases in λ_t and thus in σ_t^2 . Going further, $\lambda_t = \hat{\lambda}_{t-1}\nu$ (weakly) increases in σ_{t-1}^2 and hence $\hat{\lambda}_t$ (weakly) increases in σ_{t-1}^2 . Iterating the same argument, we conclude that, for any $t \geq 1$, $\hat{\lambda}_t$ (weakly) increases in σ_1^2 . By an argument similar to that in the proof Proposition 4, this property is preserved in the limit. \square

Proof of Proposition 5. As in Section 3, we drop all subscripts referring to the period t .

Preliminaries. We observe that $\hat{\lambda}$ is almost everywhere differentiable in σ^2 , ν^2 , and α . It is easy to check that whenever $\hat{\lambda} \in (0, 1)$, it is strictly monotone in all three parameters: (FD) is strictly increasing in α and σ^2 , and it is strictly decreasing in ν^2 , so the claim follows from Theorem 10.6 in Sundaram (1996). Moreover, fixing the other two parameters, we have $\hat{\lambda} > 0$ if and only if α is large enough, σ^2 is large enough, or ν^2 is small enough. Hence, $\hat{\lambda}$ is monotone in all three parameters. And a real-valued, monotone function is almost everywhere differentiable, which was to be proven. Going forward, we drop the qualifier “almost everywhere” when taking partial derivatives.

The agent’s ex-ante value function is given by

$$\mathcal{U}(\sigma^2, \nu^2; \alpha, \gamma) = \alpha\gamma(\mu + b(\hat{\lambda})) - \frac{\gamma}{2}\hat{\lambda}\nu^2 + 2\alpha\mu - \frac{1}{2}\sigma^2 - \frac{1}{2}(1 + (1 - \hat{\lambda})^2)b(\hat{\lambda})^2 - \frac{1}{2}(\lambda - \hat{\lambda})^2(\sigma^2 + \nu^2) - \frac{1}{2}\lambda\nu^2.$$

Part I. By the envelope theorem, we have $\frac{\partial \mathcal{U}}{\partial \hat{\lambda}} \frac{\partial \hat{\lambda}}{\partial \sigma^2} = 0$. With this, we first observe that

$$\begin{aligned} \frac{\partial \mathcal{U}}{\partial \sigma^2} &= \frac{\partial \mathcal{U}}{\partial \hat{\lambda}} \frac{\partial \hat{\lambda}}{\partial \sigma^2} - \frac{1}{2} - \frac{1}{2}(\lambda - \hat{\lambda})^2 - \frac{\partial \lambda}{\partial \sigma^2}(\lambda - \hat{\lambda})(\sigma^2 + \nu^2) - \frac{1}{2} \frac{\partial \lambda}{\partial \sigma^2} \nu^2 \\ &= -\frac{1}{2} - \frac{1}{2}(\lambda - \hat{\lambda})^2 - (1 - \lambda)(\lambda - \hat{\lambda}) - \frac{1}{2}(1 - \lambda)^2 \\ &= -\frac{1}{2}(1 + (1 - \lambda)^2 + (\lambda - \hat{\lambda})(2 - \lambda - \hat{\lambda})), \end{aligned}$$

where the second equality holds follow from $\frac{\partial \lambda}{\partial \sigma^2} = \frac{1 - \lambda}{\sigma^2 + \nu^2}$. Moreover, we have

$$\frac{\partial^2 \mathcal{U}}{\partial \hat{\lambda} \partial \sigma^2} = \frac{1}{2}(2 - \lambda - \hat{\lambda} + \lambda - \hat{\lambda}) = 1 - \hat{\lambda} > 0,$$

which in turn implies

$$\frac{\partial \mathcal{U}}{\partial \sigma^2} \leq \left. \frac{\partial \mathcal{U}}{\partial \sigma^2} \right|_{\hat{\lambda}=1} = -\frac{1}{2}.$$

Second, we observe that

$$\begin{aligned}\frac{\partial \mathcal{U}}{\partial \nu^2} &= \frac{\partial \mathcal{U}}{\partial \hat{\lambda}} \frac{\partial \hat{\lambda}}{\partial \nu^2} - \frac{\gamma}{2} \hat{\lambda} - \frac{1}{2} (\lambda - \hat{\lambda})^2 - \frac{\partial \lambda}{\partial \nu^2} (\lambda - \hat{\lambda}) (\sigma^2 + \nu^2) + \frac{1}{2} \frac{\partial \lambda}{\partial \nu^2} \sigma^2 \\ &= -\frac{\gamma}{2} \hat{\lambda} - \frac{1}{2} (\lambda - \hat{\lambda})^2 + \lambda (\lambda - \hat{\lambda}) - \frac{1}{2} \lambda^2 \\ &= -\frac{1}{2} \hat{\lambda} (\gamma + \hat{\lambda}) \leq 0,\end{aligned}$$

where the second equality holds by the envelope theorem, which gives $\frac{\partial \mathcal{U}}{\partial \hat{\lambda}} \frac{\partial \hat{\lambda}}{\partial \nu^2} = 0$, and $\frac{\partial \lambda}{\partial \sigma^2} = \frac{-\lambda}{\sigma^2 + \nu^2}$.

Part II. We first observe that

$$\frac{\partial^2 \mathcal{U}}{\partial \alpha \partial \sigma^2} = \frac{\partial^2 \mathcal{U}}{\partial \hat{\lambda} \partial \sigma^2} \frac{\partial \hat{\lambda}}{\partial \alpha}.$$

By Part I, $\frac{\partial^2 \mathcal{U}}{\partial \hat{\lambda} \partial \sigma^2} > 0$, and by our preliminary considerations, $\frac{\partial \hat{\lambda}}{\partial \alpha} \geq 0$. Hence, $\frac{\partial^2 \mathcal{U}}{\partial \alpha \partial \sigma^2} \geq 0$.

Second, we observe that

$$\frac{\partial^2 \mathcal{U}}{\partial \alpha \partial \nu^2} = \frac{\partial^2 \mathcal{U}}{\partial \hat{\lambda} \partial \nu^2} \frac{\partial \hat{\lambda}}{\partial \alpha} = -\frac{1}{2} (\gamma + 2\hat{\lambda}) \frac{\partial \hat{\lambda}}{\partial \alpha},$$

which is weakly negative again by our preliminary considerations.

Part III. By Part II, we know that

$$\frac{\partial}{\partial \alpha} \left(\frac{\partial \mathcal{U}}{\partial \sigma^2} - \frac{\partial \mathcal{U}}{\partial \nu^2} \right) \geq 0.$$

Moreover, it is easy to check that for any $\hat{\lambda} < 1$, (FD) in the proof of Proposition 2 diverges to infinity as $\alpha \rightarrow \infty$. Hence, we must have $\lim_{\alpha \rightarrow \infty} \hat{\lambda} = 1$, which in turn implies that

$$\lim_{\alpha \rightarrow \infty} \left(\frac{\partial \mathcal{U}}{\partial \sigma^2} - \frac{\partial \mathcal{U}}{\partial \nu^2} \right) = \lim_{\hat{\lambda} \rightarrow 1} \left(\frac{\partial \mathcal{U}}{\partial \sigma^2} - \frac{\partial \mathcal{U}}{\partial \nu^2} \right) = -\frac{1}{2} + \frac{1}{2} (1 + \gamma) = \frac{\gamma}{2} > 0.$$

Hence, for any $\gamma > 0$, there exists an $\check{\alpha} \in \mathbb{R}_{>0}$, so that for all $\alpha > \check{\alpha}$ and all $\sigma^2 \in \mathbb{R}_{\geq 0}$ and $\nu^2 \in \mathbb{R}_{>0}$,

$$\frac{\partial}{\partial \sigma^2} \mathcal{U}(\sigma^2, \nu^2; \alpha, \gamma) > \frac{\partial}{\partial \nu^2} \mathcal{U}(\sigma^2, \nu^2; \alpha, \gamma). \quad \square$$

Proof of Proposition 6. Since the agent thinks ahead only one period, upon delaying to receive the additional signal in period t , she expects to get it in period $t + 1$. We start by calculating the value of receiving the signal in period t , which we denote by \mathcal{U}_{now} , and the value of receiving it in period $t + 1$, which we denote by $\mathcal{U}_{\text{later}}$, for an arbitrary original variance σ_t^2 . For that, we define as $\mathcal{U}_{\text{base}}$

the agent's utility absent the additional signal, and we make two preliminary observations. First, we observe that getting the signal in t reduces the original variance from σ_t^2 to

$$\frac{\sigma_t^2 \frac{\nu^2}{d}}{\sigma_t^2 + \frac{\nu^2}{d}} = \frac{\sigma_t^2 \nu^2}{d\sigma_t^2 + \nu^2}.$$

Second, because all signals are normally distributed, we can think of getting the additional signal in period $t + 1$ as lowering the variance of the next regular signal from ν^2 to $\frac{\nu^2}{d+1}$.

Recycling calculations from the proof of Proposition 5, we obtain

$$\begin{aligned} \mathcal{U}_{\text{now}} &= - \int_{\frac{\sigma_t^2 \nu^2}{d\sigma_t^2 + \nu^2}}^{\sigma_t^2} \frac{\partial}{\partial \sigma^2} \mathcal{U}(\sigma^2, \nu^2; \alpha, \gamma) d\sigma^2 + \mathcal{U}_{\text{base}} \\ &= \frac{1}{2} \int_{\frac{\sigma_t^2 \nu^2}{d\sigma_t^2 + \nu^2}}^{\sigma_t^2} 1 + (1 - \lambda_t)^2 + (\lambda_t - \hat{\lambda}_t)(2 - \lambda_t - \hat{\lambda}_t) d\sigma^2 + \mathcal{U}_{\text{base}}. \end{aligned}$$

Similarly, we obtain

$$\begin{aligned} \mathcal{U}_{\text{later}} &= - \int_{\frac{\nu^2}{d+1}}^{\nu^2} \frac{\partial}{\partial \nu^2} \mathcal{U}(\sigma^2, \nu^2; \alpha, \gamma) d\nu^2 + \mathcal{U}_{\text{base}} \\ &= \frac{1}{2} \int_{\frac{\nu^2}{d+1}}^{\nu^2} \hat{\lambda}_t(\gamma + \hat{\lambda}_t) d\nu^2 + \mathcal{U}_{\text{base}}. \end{aligned}$$

Next, we derive the value of delaying the signal to the next period:

$$\mathcal{U}_{\text{later}} - \mathcal{U}_{\text{now}} = \frac{1}{2} \int_{\frac{\nu^2}{d+1}}^{\nu^2} \hat{\lambda}_t(\gamma + \hat{\lambda}_t) d\nu^2 - \frac{1}{2} \int_{\frac{\sigma_t^2 \nu^2}{d\sigma_t^2 + \nu^2}}^{\sigma_t^2} 1 + (1 - \lambda_t)^2 + (\lambda_t - \hat{\lambda}_t)(2 - \lambda_t - \hat{\lambda}_t) d\sigma^2.$$

Because both integrands on the right-hand side above are positive and can be bounded from above by a constant, by the dominated convergence theorem, we have

$$\begin{aligned} \lim_{\alpha \rightarrow \infty} (\mathcal{U}_{\text{later}} - \mathcal{U}_{\text{now}}) &= \frac{1}{2} \int_{\frac{\nu^2}{d+1}}^{\nu^2} \lim_{\alpha \rightarrow \infty} \hat{\lambda}_t(\gamma + \hat{\lambda}_t) d\nu^2 - \frac{1}{2} \int_{\frac{\sigma_t^2 \nu^2}{d\sigma_t^2 + \nu^2}}^{\sigma_t^2} \lim_{\alpha \rightarrow \infty} 1 + (1 - \lambda_t)^2 + (\lambda_t - \hat{\lambda}_t)(2 - \lambda_t - \hat{\lambda}_t) d\sigma^2 \\ &= \frac{1}{2} \int_{\frac{\nu^2}{d+1}}^{\nu^2} \gamma + 1 d\nu^2 - \frac{1}{2} \int_{\frac{\sigma_t^2 \nu^2}{d\sigma_t^2 + \nu^2}}^{\sigma_t^2} 1 d\sigma^2 \\ &= \frac{1}{2} \left((\gamma + 1) \frac{d}{d+1} \nu^2 - \frac{d\sigma_t^2}{d\sigma_t^2 + \nu^2} \sigma_t^2 \right), \end{aligned}$$

where the second equality uses that $\lim_{\alpha \rightarrow \infty} \hat{\lambda}_t = 1$. The above is strictly positive if and only if

$$\gamma > \frac{\sigma_t^2}{\nu^2} \frac{d\sigma_t^2 + \sigma_t^2}{d\sigma_t^2 + \nu^2} - 1.$$

Hence, for any $\sigma_t^2 \leq \nu^2$ and $\gamma > 0$, there exists $\alpha' = \alpha'(\sigma_t^2)$ such that for any $\alpha > \alpha'$, $\mathcal{U}_{\text{later}} > \mathcal{U}_{\text{now}}$.

To complete the proof, let $\sigma_1^2 \leq \nu^2$ and observe that for any $t \geq 2$, we have $\sigma_t^2 = \hat{\lambda}_{t-1}\nu^2 \leq \nu^2$. Next, we observe that, by Part I of Lemma 3, we have $\sigma_t^2 \in I := [\min\{\sigma_1^2, \hat{\lambda}_\infty\nu^2\}, \max\{\sigma_1^2, \hat{\lambda}_\infty\nu^2\}]$ for any t . Finally, notice that $\alpha' = \alpha'(\sigma_t^2)$ continuously changes with σ_t^2 . Hence, since I is compact, $\hat{\alpha} := \max_{\sigma_t^2 \in I} \alpha'(\sigma_t^2)$ exists. Fixing $\gamma > 0$, any agent with $\alpha > \hat{\alpha}$ thus delays the signal eternally. \square

References

- Akerlof, George A and William T Dickens**, “The economic consequences of cognitive dissonance,” *American Economic Review*, 1982, 72 (3), 307–319.
- Allcott, Hunt, Levi Boxell, Jacob Conway, Matthew Gentzkow, Michael Thaler, and David Yang**, “Polarization and public health: Partisan differences in social distancing during the coronavirus pandemic,” *Journal of Public Economics*, 2020, 191, 104254.
- Ba, Cuimin**, “Robust Misspecified Models and Paradigm Shifts,” https://cuiminba.com/uploads/Robust_models.pdf 2024.
- Bénabou, Roland and Jean Tirole**, “Self-Confidence and Personal Motivation,” *Quarterly Journal of Economics*, 2002, pp. 871–915.
- Benjamin, Daniel J**, “Errors in probabilistic reasoning and judgment biases,” *Handbook of Behavioral Economics: Applications and Foundations 1*, 2019, 2, 69–186.
- Benoît, Jean-Pierre, Juan Dubra, and Don A Moore**, “Does the better-than-average effect show that people are overconfident?: Two experiments,” *Journal of the European Economic Association*, 2015, 13 (2), 293–329.
- Berglas, Steven**, “How to keep A players productive,” *Harvard Business Review*, 2006, 84 (9), 104.
- Berk, Robert H**, “Limiting behavior of posterior distributions when the model is incorrect,” *The Annals of Mathematical Statistics*, 1966, 37 (1), 51–58.

- Brunnermeier, Markus K and Jonathan A Parker**, “Optimal expectations,” *American Economic Review*, 2005, *95* (4), 1092–1118.
- , **Filippos Papakonstantinou, and Jonathan A Parker**, “Optimal time-inconsistent beliefs: Misplanning, procrastination, and commitment,” *Management Science*, 2017, *63* (5), 1318–1340.
- Camerer, Colin and Dan Lovallo**, “Overconfidence and excess entry: An experimental approach,” *American Economic Review*, 1999, *89* (1), 306–318.
- Caplin, Andrew and John Leahy**, “Psychological expected utility theory and anticipatory feelings,” *Quarterly Journal of Economics*, 2001, *116* (1), 55–79.
- and **John V Leahy**, “Wishful Thinking,” https://www.nber.org/system/files/working_papers/w25707/w25707.pdf 2019.
- Charness, Gary, Aldo Rustichini, and Jeroen Van de Ven**, “Self-confidence and strategic behavior,” *Experimental Economics*, 2018, *21*, 72–98.
- Drobner, Christoph**, “Motivated beliefs and anticipation of uncertainty resolution,” *American Economic Review: Insights*, 2022, *4* (1), 89–105.
- Eil, David and Justin M. Rao**, “The Good News-Bad News Effect: Asymmetric Processing of Objective Information about Yourself,” *American Economic Journal: Microeconomics*, May 2011, *3* (2), 114–138.
- Ferrari, Joseph R and Ted Thompson**, “Impostor fears: Links with self-presentational concerns and self-handicapping behaviours,” *Personality and Individual Differences*, 2006, *40* (2), 341–352.
- Fudenberg, Drew, Giacomo Lanzani, and Philipp Strack**, “Pathwise concentration bounds for Bayesian beliefs,” *Theoretical Economics*, 2023, *18* (4), 1585–1622.
- Gagnon-Bartsch, Tristan, Matthew Rabin, and Joshua Schwartzstein**, “Channeled attention and stable errors,” <https://gagnon-bartsch.com/files/ChanneledAttention2023.pdf> 2023.

- Golman, Russell, David Hagmann, and George Loewenstein**, “Information avoidance,” *Journal of Economic Literature*, 2017, 55 (1), 96–135.
- Gottlieb, Daniel**, “Will you never learn? self deception and biases in information processing,” <https://personal.lse.ac.uk/GOTTLIED/publications/will-you-never-learn.pdf> 2014.
- Heidhues, Paul, Botond Köszegi, and Philipp Strack**, “Convergence in models of misspecified learning,” *Theoretical Economics*, 2021, 16 (1), 73–99.
- Hoffman, Mitchell and Stephen V Burks**, “Worker overconfidence: Field evidence and implications for employee turnover and firm profits,” *Quantitative Economics*, 2020, 11 (1), 315–348.
- Jevons, Herbert Stanley**, *Essays on economics*, Macmillan and Company, limited, 1905.
- Köszegi, Botond**, “Ego utility, overconfidence, and task choice,” *Journal of the European Economic Association*, 2006, 4 (4), 673–707.
- Köszegi, Botond, George Loewenstein, and Takeshi Murooka**, “Fragile self-esteem,” *Review of Economic Studies*, 2022, 89 (4), 2026–2060.
- Landier, Augustin**, “Wishful thinking: A model of optimal reality denial,” <https://pages.stern.nyu.edu/~alandier/pdfs/wishfulthinking.pdf> 2000.
- Loewenstein, George**, “Anticipation and the valuation of delayed consumption,” *The Economic Journal*, 1987, 97 (387), 666–684.
- Malmendier, Ulrike and Geoffrey Tate**, “CEO overconfidence and corporate investment,” *Journal of Finance*, 2005, 60 (6), 2661–2700.
- Mayraz, Guy**, “Priors and Desires - A Bayesian Model of Wishful Thinking and Cognitive Dissonance,” <https://mayraz.com/papers/PriorsAndDesires.pdf> 2019. Working Paper.
- Möbius, Markus M, Muriel Niederle, Paul Niehaus, and Tanya S Rosenblat**, “Managing self-confidence: Theory and experimental evidence,” *Management Science*, 2022, 68 (11), 7793–7817.

- Oster, Emily, Ira Shoulson, and E. Ray Dorsey**, “Optimal Expectations and Limited Medical Testing: Evidence from Huntington Disease,” *American Economic Review*, April 2013, *103* (2), 804–830.
- Rabin, Matthew and Joel L Schrag**, “First impressions matter: A model of confirmatory bias,” *Quarterly Journal of Economics*, 1999, *114* (1), 37–82.
- Spiegler, Ran**, “On two points of view regarding revealed preference and behavioral economics,” *The foundations of positive and normative economics: A handbook*, 2008, pp. 95–115.
- , “Behavioral economics and the atheoretical style,” *American Economic Journal: Microeconomics*, 2019, *11* (2), 173–194.
- Sundaram, Rangarajan K**, *A first course in optimization theory*, Cambridge university press, 1996.
- Yariv, Leeat**, “I’ll See it When I Believe it? A Simple Model of Cognitive Consistency,” <https://www.lyariv.com/papers/Believe.pdf> 2005.
- Zimmermann, Florian**, “The dynamics of motivated beliefs,” *American Economic Review*, 2020, *110* (2), 337–61.

Online Appendix to The Dynamics of Chosen Beliefs

Marc Kaufmann

Mats Köster

November 14, 2024

Abstract

This online appendix contains additional derivations used in the proofs provided in the main text as well as some additional results on model variations and robustness.

Keywords: Belief-Based Utility, Motivated Reasoning, Dogmas, Persistent Insecurity.

A Preliminaries: Deriving the Objective Function

A.1 The Case of $T = 2$

As in the main text, we drop all subscripts referring to the period t , and we normalize $\delta = 1$ and $\phi = 1$. Fixing a chosen mean $\hat{\mu}$ and weight on the signal $\hat{\lambda}$, the agent correctly predicts to choose

$$\tilde{a}_2(s) = (1 - \hat{\lambda})\hat{\mu} + \hat{\lambda}s$$

in $t = 2$. Using this, we can calculate the agent's anticipatory utility

$$\begin{aligned}\mathbb{E}_{\hat{\mu}, \hat{\sigma}^2}^{\theta, s} [u(\theta, \tilde{a}_2(s))] &= \alpha \mathbb{E}_{\hat{\mu}, \hat{\sigma}^2}^{\theta, s} [\theta] - \frac{1}{2} \mathbb{E}_{\hat{\mu}, \hat{\sigma}^2}^{\theta, s} [((1 - \hat{\lambda})\hat{\mu} + \hat{\lambda}s - \theta)^2] \\ &= \alpha \hat{\mu} - \frac{1}{2} \hat{\lambda} \nu^2 \\ &= \alpha(\mu + \hat{b}) - \frac{1}{2} \hat{\lambda} \nu^2,\end{aligned}$$

where \hat{b} is the chosen bias in the agent's mean belief in $t = 1$.

The agent's expected consumption utility in the first period then is

$$\begin{aligned}
\mathbb{E}_{\mu,\sigma^2}^\theta [u(\theta, \hat{a}_1)] &= \alpha\mu - \frac{1}{2}\mathbb{E}_{\mu,\sigma^2}^\theta [(\mu + \hat{b} - \theta)^2] \\
&= \alpha\mu - \frac{1}{2}\mathbb{E}_{\mu,\sigma^2}^\theta [(\mu - \theta)^2 + 2\hat{b}(\mu - \theta) + \hat{b}^2] \\
&= \alpha\mu - \frac{1}{2}\sigma^2 - \frac{1}{2}\hat{b}^2,
\end{aligned}$$

while her expected consumption utility in the second period is

$$\begin{aligned}
\mathbb{E}_{\mu,\sigma^2}^{\theta,s} [u(\theta, \hat{a}_2(s))] &= \alpha\mu - \frac{1}{2}\mathbb{E}_{\mu,\sigma^2}^s \left[\mathbb{E}_{\mu,\sigma^2}^\theta [((1 - \hat{\lambda})(\mu + \hat{b}) + \hat{\lambda}s - \theta)^2 | s] \right] \\
&= \alpha\mu - \frac{1}{2}\mathbb{E}_{\mu,\sigma^2}^s \left[\mathbb{E}_{\mu,\sigma^2}^\theta [((1 - \hat{\lambda})\mu + \hat{\lambda}s - \theta)^2 | s] \right] \\
&\quad - \mathbb{E}_{\mu,\sigma^2}^s \left[\mathbb{E}_{\mu,\sigma^2}^\theta [((1 - \hat{\lambda})\mu + \hat{\lambda}s - \theta)(1 - \hat{\lambda})\hat{b} | s] \right] - \frac{1}{2}(1 - \hat{\lambda})^2\hat{b}^2.
\end{aligned}$$

First, because $\mathbb{E}_{\mu,\sigma^2}^\theta [\theta | s] = (1 - \lambda)\mu + \lambda s$ and $\mathbb{E}_{\mu,\sigma^2}^s [s] = \mu$, we have

$$\mathbb{E}_{\mu,\sigma^2}^s \left[\mathbb{E}_{\mu,\sigma^2}^\theta [(1 - \hat{\lambda})\mu + \hat{\lambda}s - \theta | s] \right] = \mathbb{E}_{\mu,\sigma^2}^s [(\lambda - \hat{\lambda})(\mu - s)] = 0.$$

Second, because $\nu^2 = \mathbb{E}_{\mu,\sigma^2}^s [s^2 | \theta] - \theta^2$ and $\sigma^2 = \mathbb{E}_{\mu,\sigma^2}^\theta [\theta^2] - \mu^2$,

$$\mathbb{E}_{\mu,\sigma^2}^s [(\mu - s)^2] = \mathbb{E}_{\mu,\sigma^2}^s [s^2] - \mu^2 = \mathbb{E}_{\mu,\sigma^2}^\theta [\mathbb{E}_{\mu,\sigma^2}^s [s^2 | \theta]] - \mu^2 = \mathbb{E}_{\mu,\sigma^2}^\theta [\nu^2 + \theta^2] - \mu^2 = \nu^2 + \sigma^2,$$

where the second equality holds by the law of iterated expectations.

Third, since $\mathbb{E}_{\mu,\sigma^2}^\theta [\theta | s] = (1 - \lambda)\mu + \lambda s$ and $\mathbb{E}_{\mu,\sigma^2}^s [(\mu - s)^2] = \sigma^2 + \nu^2$,

$$\begin{aligned}
&\mathbb{E}_{\mu,\sigma^2}^s \left[\mathbb{E}_{\mu,\sigma^2}^\theta [((1 - \hat{\lambda})\mu + \hat{\lambda}s - \theta)^2 | s] \right] \\
&= \mathbb{E}_{\mu,\sigma^2}^s \left[\mathbb{E}_{\mu,\sigma^2}^\theta [((\lambda - \hat{\lambda})(\mu - s) + (1 - \lambda)\mu + \lambda s - \theta)^2 | s] \right] \\
&= (\lambda - \hat{\lambda})^2 \mathbb{E}_{\mu,\sigma^2}^s [(\mu - s)^2] + \mathbb{E}_{\mu,\sigma^2}^s \left[\mathbb{E}_{\mu,\sigma^2}^\theta [((1 - \lambda)\mu + \lambda s - \theta)^2 | s] \right] \\
&= (\lambda - \hat{\lambda})^2 (\sigma^2 + \nu^2) + \lambda \nu^2.
\end{aligned}$$

We thus conclude that

$$\mathbb{E}_{\mu,\sigma^2}^{\theta,s} [u(\theta, \hat{a}_2(s))] = \alpha\mu - \frac{1}{2}(\lambda - \hat{\lambda})^2 (\sigma^2 + \nu^2) - \frac{1}{2}\lambda\nu^2 - \frac{1}{2}(1 - \hat{\lambda})^2\hat{b}^2.$$

Putting all of this together, and plugging in the optimal mean bias \hat{b} derived in Proposition 1, we arrive at the objective function in Eq. (5) of the main text.

A.2 The General Case

Consider period t , and fix $\hat{\mu}_t$ as well as $\hat{\lambda}_t$. For every $\tau \in \{t+1, \dots, T-1\}$, let

$$\tilde{\lambda}_\tau = \frac{\hat{\lambda}_t}{1 + (\tau - t)\hat{\lambda}_t}$$

be the weight on signal s_τ implied by Bayesian updating based on the chosen weight on the next signal, $\hat{\lambda}_t$. We first observe that, for every $\tau \in \{t+1, \dots, T-1\}$, we have

$$\prod_{\ell=t+1}^{\tau-1} 1 - \tilde{\lambda}_\ell = \prod_{\ell=t+1}^{\tau-1} \frac{1 + (\ell - (t+1))\hat{\lambda}_t}{1 + (\ell - t)\hat{\lambda}_t} = \frac{1}{1 + \hat{\lambda}_t} \frac{1 + \hat{\lambda}_t}{1 + 2\hat{\lambda}_t} \dots \frac{1 + (\tau - (t+2))\hat{\lambda}_t}{1 + (\tau - (t+1))\hat{\lambda}_t} = \frac{1}{1 + (\tau - (t+1))\hat{\lambda}_t}$$

and, for every $\ell \in \{t+1, \dots, \tau-1\}$,

$$\tilde{\lambda}_\ell \prod_{k=\ell+1}^{\tau-1} 1 - \tilde{\lambda}_k = \frac{\hat{\lambda}_t}{1 + (\ell - t)\hat{\lambda}_t} \frac{1 + (\ell - t)\hat{\lambda}_t}{1 + (\tau - (t+1))\hat{\lambda}_t} = \frac{\hat{\lambda}_t}{1 + (\tau - (t+1))\hat{\lambda}_t}.$$

Thus, from the perspective of period t , the agent expects to choose

$$\begin{aligned} \tilde{a}_\tau(\mathbf{s}_t^\tau) &= \hat{\mu}_t(1 - \hat{\lambda}_t) \prod_{\ell=t+1}^{\tau-1} (1 - \tilde{\lambda}_\ell) + s_t \hat{\lambda}_t \prod_{k=t+1}^{\tau-1} (1 - \tilde{\lambda}_k) + \sum_{\ell=t+1}^{\tau-1} s_\ell \tilde{\lambda}_\ell \prod_{k=\ell+1}^{\tau-1} (1 - \tilde{\lambda}_k) \\ &= \hat{\mu}_t \frac{1 - \hat{\lambda}_t}{1 + (\tau - (t+1))\hat{\lambda}_t} + \underbrace{\frac{\hat{\lambda}_t(\tau - t)}{1 + (\tau - (t+1))\hat{\lambda}_t}}_{=: w_t^\tau} \underbrace{\frac{1}{(\tau - t)} \sum_{\ell=t}^{\tau-1} s_\ell}_{=: \mathcal{S}_t^\tau} \\ &= (1 - w_t^\tau(\hat{\lambda}_t))(\mu_t + \hat{b}_t) + w_t^\tau(\hat{\lambda}_t)\mathcal{S}_t^\tau, \end{aligned}$$

where $w_t^\tau \in [0, 1]$ and $\mathcal{S}_t^\tau \sim \mathcal{N}(\theta, \nu^2/(\tau-t))$.

By the same arguments as in Online Appendix A.1, for every $\tau \in \{t+1, \dots, T\}$, we have

$$\mathbb{E}_{\mu_t, \sigma_t^2}^{\theta, \mathcal{S}_t^\tau} [u(\theta, \tilde{a}_\tau(\mathbf{s}_t^\tau))] = \alpha \mu_t - \frac{1}{2} (w_t^\tau(\lambda_t) - w_t^\tau(\hat{\lambda}_t))^2 \left(\sigma_t^2 + \frac{\nu^2}{\tau - t} \right) - \frac{1}{2} w_t^\tau(\lambda_t) \nu^2 - \frac{1}{2} (1 - w_t^\tau(\hat{\lambda}_t))^2 \hat{b}_t^2,$$

where λ_t is the weight on signal s_t implied by the agent's genuine belief $\mathcal{N}(\mu_t, \sigma_t^2)$ and Bayes' rule.

Moreover, again by the same arguments as in Online Appendix A.1, we have

$$\mathbb{E}_{\mu_t, \sigma_t^2}^\theta [u(\theta, \hat{a}_t)] = \alpha \mu_t - \frac{1}{2} \sigma_t^2 - \frac{1}{2} \hat{b}_t^2$$

Finally, anticipatory utility from imagining the future in $\tau \in \{t+1, \dots, T\}$ periods is given by

$$\mathbb{E}_{\hat{\mu}_t, \hat{\sigma}_t^2}^{\theta, \mathbf{s}_t^\tau} \left[\phi_{\tau-t} u(\theta, \tilde{a}_\tau(\mathbf{s}_\tau^t)) \right] = \phi_{\tau-t} \left(\alpha(\mu + \hat{b}_t) - \frac{1}{2} w_t^\tau(\hat{\lambda}_t) \frac{\nu^2}{\tau-t} \right).$$

Putting all of this together, we conclude that the agent chooses \hat{b}_t and $\hat{\lambda}_t$ as to maximize

$$\begin{aligned} -\frac{1}{2} \hat{b}_t^2 - \frac{1}{2} \sum_{\tau=t+1}^T \delta_{\tau-t} \left((w_t^\tau(\lambda_t) - w_t^\tau(\hat{\lambda}_t))^2 \left(\sigma_t^2 + \frac{\nu^2}{\tau-t} \right) + (1 - w_t^\tau(\hat{\lambda}_t))^2 \hat{b}_t^2 \right) \\ + \gamma \sum_{\tau=t+1}^T \phi_{\tau-t} \left(\alpha \hat{b}_t - \frac{1}{2} w_t^\tau(\hat{\lambda}_t) \frac{\nu^2}{\tau-t} \right). \end{aligned}$$

B Information Delay and Discounting

In this online appendix, we want to argue that delaying information becomes even more attractive when the agent discounts the future less. To make this point as concisely as possible, suppose that the additional signal is fully revealing the state while the regular signals are pure noise.

When receiving the signal now, the agent has no incentive to remain uncertain, since there are no future signals. She thus chooses to be dogmatic. In expectation, receiving the signal now yields

$$\mathcal{U}_{\text{now}} = \alpha\mu - \frac{1}{2} \hat{b}_{\text{now}}^2 + \sum_{t=1}^{\infty} \delta_t \left(\alpha\mu - \frac{1}{2} \hat{b}_{\text{now}}^2 \right) + \gamma \sum_{t=1}^{\infty} \phi_t \alpha \left(\mu + \hat{b}_{\text{now}} \right).$$

As a result, the agent optimally chooses a mean bias of $\hat{b}_{\text{now}} = \frac{\alpha\gamma\Phi}{1+\Delta}$. When delaying the signal, the agent plans to receive it in the next period, and as long as she does not choose to be dogmatic, she expects to take the correct action in all future periods. The agent is therefore indifferent between choosing any non-zero variance, and delaying the signal yields an expected utility of

$$\mathcal{U}_{\text{later}} = \alpha\mu - \frac{1}{2} \sigma_1^2 - \frac{1}{2} \hat{b}_{\text{later}}^2 + \sum_{t=1}^{\infty} \delta_t \alpha\mu + \gamma \sum_{t=1}^{\infty} \phi_t \alpha \left(\mu + \hat{b}_{\text{later}} \right).$$

The optimal mean bias in this case is given by $\hat{b}_{\text{later}} = \alpha\gamma\Phi > \hat{b}_{\text{now}}$. Comparing both, we get

$$\mathcal{U}_{\text{later}} - \mathcal{U}_{\text{now}} = -\frac{1}{2} \sigma_1^2 + \Phi^2 \frac{\Delta}{1+\Delta} \frac{\alpha^2 \gamma^2}{2},$$

which increases in Φ and Δ . The agent eternally delays information if and only if $\alpha > \sqrt{\frac{\sigma_1^2(1+\Delta)}{\gamma^2 \Phi^2 \Delta}}$.

C The Cost of Belief Distortions

Following Brunnermeier and Parker (2005), we model an implicit cost of distorting one’s belief, which comes in the form of worse actions. We can, however, make this cost explicit, allowing us to compare our model more easily to alternative approaches like Caplin and Leahy (2019).

We start by re-writing our model in terms of an explicit “cost from distorted beliefs.” Consider an original belief $\theta \sim \mathcal{N}(\mu_t, \sigma_t^2)$ and a chosen belief $\theta \sim \mathcal{N}(\hat{\mu}_t, \hat{\sigma}_t^2)$. We define as

$$C((\mu_t, \sigma_t^2), (\hat{\mu}_t, \hat{\sigma}_t^2)) := \mathbb{E}_{\mu_t, \sigma_t^2} \left[u(\theta, \mu_t) + \sum_{\tau=t+1}^T \delta_{\tau-t} u(\theta, \mu_\tau(\mathbf{s}_t^\tau)) - u(\theta, \hat{\mu}_t) - \sum_{\tau=t+1}^T \delta_{\tau-t} u(\theta, \tilde{\mu}_\tau(\mathbf{s}_t^\tau)) \right]$$

the expected loss in consumption utility from choosing, and acting on, a belief different from the original one.¹ Using this new notation, the belief our agent chooses in period t maximizes

$$\underbrace{\mathbb{E}_{\hat{\mu}_t, \hat{\sigma}_t^2} \left[\sum_{\tau=t+1}^T \phi_{\tau-t} u(\theta, \tilde{\mu}_\tau(\mathbf{s}_t^\tau)) \right]}_{\text{anticipatory utility}} - \underbrace{\frac{1}{\gamma} C((\mu_t, \sigma_t^2), (\hat{\mu}_t, \hat{\sigma}_t^2))}_{\text{cost of belief distortions}}.$$

Notice that $C((\mu_t, \sigma_t^2), (\hat{\mu}_t, \hat{\sigma}_t^2)) \geq 0$, holding with equality if and only if $(\mu_t, \sigma_t^2) = (\hat{\mu}_t, \hat{\sigma}_t^2)$. We have thus re-written our model in terms of a proper cost function, as in Caplin and Leahy (2019).

We conclude this section by comparing our cost function above to the one used in Caplin and Leahy (2019, Eq. (2) and Section 5.1). Similar calculations as in Appendix A.2 yield

$$C((\mu_t, \sigma_t^2), (\hat{\mu}_t, \hat{\sigma}_t^2)) = \frac{1}{2}(\hat{\mu}_t - \mu_t)^2 + \frac{1}{2} \sum_{\tau=t+1}^T \delta_{\tau-t} \left((w_t^\tau(\lambda_t) - w_t^\tau(\hat{\lambda}_t))^2 \left(\sigma_t^2 + \frac{\nu^2}{\tau-t} \right) + (1 - w_t^\tau(\hat{\lambda}_t))^2 (\hat{\mu}_t - \mu_t)^2 \right).$$

Caplin and Leahy (2019) define costs of belief distortions via the Kullback-Leibler (KL) divergence from the original to the chosen belief. Setting $\Delta_\tau := 1 + \sum_{\tau=t+1}^T \delta_{\tau-t}$, their cost function is

$$C_{\text{CL}}((\mu_t, \sigma_t^2), (\hat{\mu}_t, \hat{\sigma}_t^2)) := \frac{1}{\Delta_\tau} \left(\text{KL}((\hat{\mu}_t, \hat{\sigma}_t^2) \| (\mu_t, \sigma_t^2)) + \sum_{\tau=t+1}^T \delta_{\tau-t} \text{KL}((\tilde{\mu}_\tau(\mathbf{s}_t^\tau), \hat{\sigma}_\tau^2) \| (\mu_\tau(\mathbf{s}_t^\tau), \sigma_\tau^2)) \right).$$

Notice that for normal distributions $\mathcal{N}(\mu_t, \sigma_t^2)$ and $\mathcal{N}(\hat{\mu}_t, \hat{\sigma}_t^2)$, we have²

$$\text{KL}((\hat{\mu}_t, \hat{\sigma}_t^2) \| (\mu_t, \sigma_t^2)) = \log \left(\frac{\sigma_t}{\hat{\sigma}_t} \right) + \frac{\hat{\sigma}_t^2 + (\hat{\mu}_t - \mu_t)^2}{2\sigma_t^2} - \frac{1}{2}.$$

¹ For a vector of signal realizations $\mathbf{s}_t^\tau = (s_t, \dots, s_{\tau-1})$, we denote a Bayesian’s posterior mean in τ as $\mu_t(\mathbf{s}_t^\tau)$.

² See https://en.wikipedia.org/wiki/KullbackLeibler_divergence (accessed on November 12, 2024).

Again using similar calculations as in Appendix A.2 (and setting $\delta_0 = 1$), we obtain

$$C_{\text{CL}}((\mu_t, \sigma_t^2), (\hat{\mu}_t, \hat{\sigma}_t^2)) = \frac{1}{\Delta_\tau \sigma_t^2} C((\mu_t, \sigma_t^2), (\hat{\mu}_t, \hat{\sigma}_t^2)) + \sum_{\tau=t}^T \frac{\delta_{\tau-t}}{\Delta_\tau} \left\{ \log \left(\frac{\sigma_\tau}{\hat{\sigma}_\tau} \right) + \frac{1}{2} \frac{\hat{\sigma}_\tau^2}{\sigma_\tau^2} - \frac{1}{2} \right\}.$$

We observe that the cost function in Caplin and Leahy (2019) places a higher cost on distortions (in either direction) of the variance. In particular, for any fixed μ_t , σ_t^2 , and $\hat{\mu}_t$, $C_{\text{CL}}((\mu_t, \sigma_t^2), (\hat{\mu}_t, \hat{\sigma}_t^2))$ approaches infinity as $\hat{\sigma}_t^2 \rightarrow 0$. Hence, with this alternative cost function, the agent would not become dogmatic. Otherwise, their cost function has a shape similar to our utility-based one.

D Alternative Models

D.1 Distorting the Signal Structure

We study an alternative model in which the agent distorts her beliefs about the (distribution) of the signals. In doing so, we focus on short-run beliefs. The signal is drawn from $\mathcal{N}(\theta, \nu^2)$. Before observing the signal, the agent can choose a belief about the signal structure $s \sim \mathcal{N}(\hat{\mu}_s, \hat{\nu}^2)$. This is equivalent to choosing a bias $\hat{b} \in \mathbb{R}$ such that $\hat{\mu}_s = \theta + \hat{b}$ and a weight on the signal $\hat{\lambda}_s \in [0, 1]$. The agent starts out with a prior $\theta \sim \mathcal{N}(\mu_1, \sigma_1^2)$, and updates via Bayes' rule using the chosen signal structure. In the second period, the agent then has to act on the belief implied by her chosen signal structure. We analyze two versions of the model that differ in when exactly the agent chooses her belief. To make the results comparable to those in Section 3, we again normalize $\delta_1 = 1$ and $\phi_1 = 1$.

Ex ante choice Suppose that the agent chooses her belief about the signal structure *before* observing the signal realization. The exact timing is summarized in Figure 1.




Figure 1: *Timing of events with ex ante choice of signal structure.*

We start by observing that, for a given \hat{b} and $\hat{\lambda}_s$, the agent correctly predicts to choose

$$a_2(s) = \mathbb{E}_{\mu_1, \sigma_1^2, \hat{b}, \hat{\nu}^2}^\theta [\theta | s] = (1 - \hat{\lambda}_s)\mu_1 + \hat{\lambda}_s(s - \hat{b})$$

in the second period. Importantly, although the agent biases her belief about the signal structure, her (posterior) belief about the state satisfies the law of iterated expectations:

$$\mathbb{E}_{\mu_1, \sigma_1^2, \hat{b}, \hat{\nu}^2}^s \left[\mathbb{E}_{\mu_1, \sigma_1^2, \hat{b}, \hat{\nu}^2}^\theta [\theta | s] \right] = \mu_1.$$

As a consequence, the agent cannot be systematically overoptimistic about the state.

Next, we observe that in $t = 1$, the agent chooses $a_1 = \mu_1$, resulting in an expected loss of σ_1^2 . Combining both observations, we conclude that the agent chooses \hat{b} and $\hat{\lambda}_s$ as to maximize

$$\underbrace{\alpha\mu_1 - \frac{1}{2}\sigma_1^2 + \alpha\mu_1 - \frac{1}{2}\mathbb{E}_{\mu_1, \sigma_1^2, \theta, \nu^2}^{\theta, s} \left[((1 - \hat{\lambda}_s)\mu_1 + \hat{\lambda}_s(s - \hat{b}) - \theta)^2 \right]}_{\text{expected consumption utility}} + \underbrace{\gamma \left(\alpha\mu_1 - \frac{1}{2}(1 - \hat{\lambda}_s)\sigma_1^2 \right)}_{\text{anticipatory utility}}.$$

Because the agent's belief satisfies the law of iterated expectations, choosing $\hat{b} \neq 0$ has no benefit in terms of anticipatory utility. Since $\hat{b} \neq 0$ results in worse actions, however, the agent optimally chooses $\hat{b} = 0$. This, in turn, implies that the agent chooses $\hat{\lambda}_s$ as to maximize

$$-\frac{1}{2}\mathbb{E}_{\mu_1, \sigma_1^2, \theta, \nu^2}^{\theta, s} \left[((1 - \hat{\lambda}_s)\mu_1 + \hat{\lambda}_s s - \theta)^2 \right] - \frac{\gamma}{2}(1 - \hat{\lambda}_s)\sigma_1^2.$$

By the same arguments as in Appendix A.1, the above can be re-written as

$$-\frac{1}{2}(\lambda - \hat{\lambda}_s)^2(\sigma_1^2 + \nu^2) - \frac{\gamma}{2}(1 - \hat{\lambda}_s)\sigma_1^2.$$

This objective function is strictly concave in $\hat{\lambda}_s$, so that the optimal weight on the signal satisfies

$$(\lambda - \hat{\lambda}_s)(\sigma_1^2 + \nu^2) + \frac{\gamma}{2}\sigma_1^2 \geq 0.$$

We conclude:

Proposition 1 (Ex-Ante Optimal Signal Structure).

The agent chooses $\hat{\mu}_s = \theta$ and $\hat{\lambda}_s = \min \{1, \lambda \frac{\gamma+2}{2}\}$, which results in a posterior mean of

$$\mu_2(s) = \begin{cases} \mu_1 + \lambda \frac{\gamma+2}{2}(s - \mu_1) & \text{if } \lambda \leq \frac{2}{2+\gamma}, \\ s & \text{if } \lambda > \frac{2}{2+\gamma}. \end{cases}$$

Contrasting this result with Propositions 1 and 2 in the main text, there are three stark differences. First, an agent who chooses a belief about the signal structure, is, on average, well calibrated. Second, the agent’s posterior is independent of α . Third, the agent is necessarily erratic, as overstating the precision of the signal is the only way in which she can reduce her anxiety regarding the second-period action. These results seem less in line with anecdotal evidence.

Ex post choice Suppose that the agent chooses her belief about the signal structure *after* observing the signal realization. To make the results comparable to the ones above (and those in Section 3), we assume that the agent observes the signal in $t = 1$; before she feels anticipatory utility, but after she chose her first-period action. The exact timing is summarized in Figure 2.




Figure 2: *Timing of events with ex post choice of signal structure.*

Fix a signal realization s . For a given \hat{b} and $\hat{\lambda}_s$, in $t = 2$, the agent predicts to choose

$$a_2(s) = (1 - \hat{\lambda}_s)\mu_1 + \hat{\lambda}_s(s - \hat{b}).$$

Hence, the agent chooses \hat{b} and $\hat{\lambda}_s$ as to maximize

$$-\frac{1}{2}\mathbb{E}_{\mu_1, \sigma_1^2}^{\theta} \left[((1 - \hat{\lambda}_s)\mu_1 + \hat{\lambda}_s(s - \hat{b}) - \theta)^2 \right] + \gamma \left(\alpha((1 - \hat{\lambda}_s)\mu_1 + \hat{\lambda}_s(s - \hat{b})) - \frac{1}{2}(1 - \hat{\lambda}_s)\sigma_1^2 \right).$$

For a given $\hat{\lambda}_s$, the objective function is strictly concave in \hat{b} . The optimal \hat{b} thus satisfies

$$\hat{\lambda}_s((1 - \hat{\lambda}_s)\mu_1 + \hat{\lambda}_s(s - \hat{b}) - \mu_1) - \gamma\alpha\hat{\lambda}_s = 0,$$

or, equivalently

$$-\hat{b} = \mu_1 - s + \frac{\gamma\alpha}{\hat{\lambda}_s}. \tag{1}$$

Using (1), the agent chooses $\hat{\lambda}_s$ as to maximize

$$-\frac{1}{2}\mathbb{E}_{\mu_1, \sigma_1^2}^{\theta} \left[(\mu_1 + \alpha\gamma - \theta)^2 \right] + \gamma \left(\alpha(\mu_1 + \alpha\gamma) - \frac{1}{2}(1 - \hat{\lambda}_s)\sigma_1^2 \right),$$

and thus chooses $\hat{\lambda}_s = 1$. We conclude:

Proposition 2 (Ex-Post Optimal Signal Structure).

The agent chooses $\hat{\mu}_s = \theta - (\mu_1 - s) - \gamma\alpha$ and $\hat{\lambda}_s = 1$, resulting in a posterior mean of $\mu_2 = \mu_1 + \alpha\gamma$.

Contrasting this result with Propositions 1 and 2 in the main text, we make two observations. First, when choosing her belief about the signal structure, the agent arrives at the same posterior belief as an agent with the same preferences who chooses a dogmatic prior. (Notice that, for any preference parameters, we can indeed find a signal structure that is sufficiently uninformative, such that a dogmatic prior is optimal.) Second, the posterior belief is independent of the signal. An agent who chooses her prior belief, in contrast, responds to sufficiently informative signals.

D.2 Anticipating Anticipatory Utility

In this online appendix, we consider an agent who anticipates future anticipatory utility (as in Brunnermeier and Parker, 2005, Brunnermeier et al., 2017). Notice that our analysis of short-run beliefs in Section 3 still applies. We thus focus on long-run beliefs (i.e., $T = \infty$). We show that, irrespective of her preferences over the state, such an agent eventually develops a dogma.

Consider an arbitrary period $t \geq 1$. We start by calculating the anticipatory utility that the agent expects to feel in a future period $\tau \geq t + 1$. When choosing $\hat{\mu}_t \in \mathbb{R}$ and $\hat{\lambda}_t \in [0, 1]$ in period t , the agent expects to hold a period- τ belief that is normally distributed with

$$\tilde{\mu}_\tau(\mathbf{s}_t^\tau) = \hat{\mu}_t(1 - w_t^\tau(\hat{\lambda}_t)) + w_t^\tau(\hat{\lambda}_t) \frac{1}{\tau - t} \sum_{\ell=t}^{\tau-1} s_\ell \quad \text{and} \quad \tilde{\sigma}_\tau^2 = \frac{\hat{\lambda}_t}{1 + (\tau - (t + 1))\hat{\lambda}_t} \nu^2.$$

Notice that the anticipated mean belief depends on the signal realizations $\mathbf{s}_t^\tau = (s_t, \dots, s_{\tau-1})$ while the anticipated variance of her belief does not. The agent forms an expectation based on her original belief in period t , and her expected mean belief in period τ is thus given by

$$\mathbb{E}_{\mu_t, \sigma_t^2}^{\mathbf{s}_t^\tau} [\tilde{\mu}_\tau(\mathbf{s}_t^\tau)] = \mu_t + (1 - w_t^\tau(\hat{\lambda}_t)) \hat{b}_t.$$

Hence, from the perspective of period t , the expected anticipatory utility in period τ equals

$$\gamma \mathbb{E}_{\mu_t, \sigma_t^2}^{\mathbf{s}_t^\tau} \left[\sum_{\ell=\tau+1}^{\infty} \phi_{\ell-\tau} \left\{ \alpha \tilde{\mu}_\tau(\mathbf{s}_t^\tau) - \frac{1}{2} \tilde{\sigma}_\tau^2 \right\} \right] = \gamma \alpha \Phi \left(\mu_t + (1 - w_t^\tau(\hat{\lambda}_t)) \hat{b}_t \right) - \frac{\gamma}{2} \sum_{\ell=\tau+1}^{\infty} \phi_{\ell-\tau} w_t^\ell(\hat{\lambda}_t) \frac{\nu^2}{\ell - t}.$$

Following Brunnermeier et al. (2017), we assume that the agent discounts anticipatory utility in $\tau - t$ periods from now with the same discount factor $\delta_{\tau-t}$ that she applies to consumption utility. Setting $\delta_0 := 1$, and recycling calculations from Appendix A.2, the agent's objective function is

$$-\frac{1}{2}\hat{b}_t^2 - \frac{1}{2} \sum_{\tau=t+1}^{\infty} \delta_{\tau-t} \left((w_t^\tau(\lambda_t) - w_t^\tau(\hat{\lambda}_t))^2 \left(\sigma_t^2 + \frac{\nu^2}{\tau-t} \right) + (1 - w_t^\tau(\hat{\lambda}_t))^2 \hat{b}_t^2 \right) + \gamma \alpha \hat{b}_t \Phi \left(1 + \sum_{\tau=t+1}^{\infty} \delta_{\tau-t} (1 - w_t^\tau(\hat{\lambda}_t)) \right) - \frac{\gamma}{2} \sum_{\tau=t+1}^{\infty} w_t^\tau(\hat{\lambda}_t) \frac{\nu^2}{\tau-t} \sum_{\ell=t+1}^{\tau} \phi_{\ell-t} \delta_{\tau-\ell}. \quad (2)$$

Proposition 3 (Anticipating Anticipatory Utility Results in Dogmatic Beliefs).

The agent develops a dogma. In particular, $\hat{\lambda}_\infty = 0$ and, for any t large enough, $\hat{\mu}_{t+1} - \hat{\mu}_t = \alpha\gamma\Phi$.

Proof. As for the agent who does not anticipate future anticipatory utility, for any fixed $\hat{\lambda}_t \in [0, 1]$, the objective function is concave in \hat{b}_t . Fixing $\hat{\lambda}_t \in [0, 1]$, we thus obtain the optimal mean bias,

$$\hat{b}_t = \alpha\gamma\Phi \frac{1 + \sum_{\tau=t+1}^{\infty} \delta_{\tau-t} (1 - w_t^\tau(\hat{\lambda}_t))}{1 + \sum_{\tau=t+1}^{\infty} \delta_{\tau-t} (1 - w_t^\tau(\hat{\lambda}_t))^2}, \quad (3)$$

by solving the first-order condition. Plugging (3) into (2), the agent chooses $\hat{\lambda}_t$ as to maximize

$$\begin{aligned} \tilde{f}(\hat{\lambda}_t) := & \frac{1}{2} \alpha^2 \gamma^2 \Phi^2 \frac{\left(1 + \sum_{\tau=t+1}^{\infty} \delta_{\tau-t} (1 - w_t^\tau(\hat{\lambda}_t)) \right)^2}{1 + \sum_{\tau=t+1}^{\infty} \delta_{\tau-t} (1 - w_t^\tau(\hat{\lambda}_t))^2} - \frac{\gamma}{2} \sum_{\tau=t+1}^{\infty} w_t^\tau(\hat{\lambda}_t) \frac{\nu^2}{\tau-t} \sum_{\ell=t+1}^{\tau} \phi_{\ell-t} \delta_{\tau-\ell} \\ & - \frac{1}{2} \sum_{\tau=t+1}^{\infty} \delta_{\tau-t} (w_t^\tau(\lambda_t) - w_t^\tau(\hat{\lambda}_t))^2 \left(\sigma_t^2 + \frac{\nu^2}{\tau-t} \right). \end{aligned}$$

Next, we show $\hat{\lambda}_t < \lambda_t$ for any $\lambda_t \in (0, 1)$. Recycling calculations from the proof of Lemma 1,

$$\begin{aligned} \tilde{f}'(\hat{\lambda}_t) = & \frac{1}{2} \alpha^2 \gamma^2 \Phi^2 \frac{\partial}{\partial \hat{\lambda}_t} \left(\frac{\left(1 + \sum_{\tau=t+1}^{\infty} \delta_{\tau-t} (1 - w_t^\tau(\hat{\lambda}_t)) \right)^2}{1 + \sum_{\tau=t+1}^{\infty} \delta_{\tau-t} (1 - w_t^\tau(\hat{\lambda}_t))^2} \right) \\ & - \frac{\gamma}{2} \nu^2 \sum_{\tau=t+1}^{\infty} \frac{1}{(1 + (\tau - (t+1))\hat{\lambda}_t)^2} \sum_{\ell=t+1}^{\tau} \phi_{\ell-t} \delta_{\tau-\ell} + \nu^2 \sum_{\tau=t+1}^{\infty} \delta_{\tau-t} \frac{(\tau-t)}{(1 + (\tau - (t+1))\hat{\lambda}_t)^3} \frac{\lambda_t - \hat{\lambda}_t}{1 - \lambda_t}. \end{aligned}$$

With the above, it is sufficient to verify that for any $\hat{\lambda}_t \in [0, 1]$, we have

$$\frac{\partial}{\partial \hat{\lambda}_t} \left(\frac{\left(1 + \sum_{\tau=t+1}^{\infty} \delta_{\tau-t} (1 - w_t^\tau(\hat{\lambda}_t)) \right)^2}{1 + \sum_{\tau=t+1}^{\infty} \delta_{\tau-t} (1 - w_t^\tau(\hat{\lambda}_t))^2} \right) \leq 0. \quad (4)$$

Notice that the inequality in (4) holds if and only if

$$\begin{aligned} & -2 \left(1 + \sum_{\tau=t+1}^{\infty} \delta_{\tau-t} (1 - w_t^\tau(\hat{\lambda}_t)) \right) \left(\sum_{\tau=t+1}^{\infty} \delta_{\tau-t} \frac{\partial}{\partial \hat{\lambda}_t} w_t^\tau(\hat{\lambda}_t) \right) \left(1 + \sum_{\tau=t+1}^{\infty} \delta_{\tau-t} (1 - w_t^\tau(\hat{\lambda}_t))^2 \right) \\ & \quad + 2 \left(\sum_{\tau=t+1}^{\infty} \delta_{\tau-t} (1 - w_t^\tau(\hat{\lambda}_t)) \frac{\partial}{\partial \hat{\lambda}_t} w_t^\tau(\hat{\lambda}_t) \right) \left(1 + \sum_{\tau=t+1}^{\infty} \delta_{\tau-t} (1 - w_t^\tau(\hat{\lambda}_t)) \right)^2 \leq 0 \end{aligned}$$

or, equivalently,

$$\begin{aligned} & - \left(\sum_{\tau=t+1}^{\infty} \delta_{\tau-t} \frac{\partial}{\partial \hat{\lambda}_t} w_t^\tau(\hat{\lambda}_t) \right) \left(1 + \sum_{\tau=t+1}^{\infty} \delta_{\tau-t} (1 - w_t^\tau(\hat{\lambda}_t))^2 \right) \\ & \quad + \left(\sum_{\tau=t+1}^{\infty} \delta_{\tau-t} (1 - w_t^\tau(\hat{\lambda}_t)) \frac{\partial}{\partial \hat{\lambda}_t} w_t^\tau(\hat{\lambda}_t) \right) \left(1 + \sum_{\tau=t+1}^{\infty} \delta_{\tau-t} (1 - w_t^\tau(\hat{\lambda}_t)) \right) \leq 0. \end{aligned}$$

Recall, again from the proof of Lemma 1, that

$$w_t^\tau(\hat{\lambda}_t) = \frac{\hat{\lambda}_t(\tau - t)}{1 + (\tau - (t + 1))\hat{\lambda}_t} \quad \text{and} \quad \frac{\partial}{\partial \hat{\lambda}_t} w_t^\tau(\hat{\lambda}_t) = \frac{(\tau - t)}{(1 + (\tau - (t + 1))\hat{\lambda}_t)^2}.$$

Hence, we can conclude that

$$\begin{aligned} & \left(\sum_{\tau=t+1}^{\infty} \delta_{\tau-t} \frac{\partial}{\partial \hat{\lambda}_t} w_t^\tau(\hat{\lambda}_t) \right) \left(\sum_{\tau=t+1}^{\infty} \delta_{\tau-t} (1 - w_t^\tau(\hat{\lambda}_t)) w_t^\tau(\hat{\lambda}_t) \right) \\ & \quad - \left(\sum_{\tau=t+1}^{\infty} \delta_{\tau-t} w_t^\tau(\hat{\lambda}_t) \frac{\partial}{\partial \hat{\lambda}_t} w_t^\tau(\hat{\lambda}_t) \right) \left(1 + \sum_{\tau=t+1}^{\infty} \delta_{\tau-t} (1 - w_t^\tau(\hat{\lambda}_t)) \right) \\ & = \hat{\lambda}_t (1 - \hat{\lambda}_t) \left(\sum_{\tau=t+1}^{\infty} \delta_{\tau-t} \frac{(\tau - t)}{(1 + (\tau - (t + 1))\hat{\lambda}_t)^2} \right)^2 \\ & \quad - \hat{\lambda}_t \left(\sum_{\tau=t+1}^{\infty} \delta_{\tau-t} \frac{(\tau - t)^2}{(1 + (\tau - (t + 1))\hat{\lambda}_t)^3} \right) \left(1 + \sum_{\tau=t+1}^{\infty} \delta_{\tau-t} \frac{1 - \hat{\lambda}_t}{1 + (\tau - (t + 1))\hat{\lambda}_t} \right) \\ & \leq \hat{\lambda}_t (1 - \hat{\lambda}_t) \left(\sum_{\tau=t+1}^{\infty} \delta_{\tau-t}^2 \frac{(\tau - t)^2}{(1 + (\tau - (t + 1))\hat{\lambda}_t)^4} \right) - \hat{\lambda}_t \left(\sum_{\tau=t+1}^{\infty} \delta_{\tau-t} \frac{(\tau - t)^2}{(1 + (\tau - (t + 1))\hat{\lambda}_t)^3} \right) \\ & \leq \hat{\lambda}_t (1 - \hat{\lambda}_t) \left(\sum_{\tau=t+1}^{\infty} \delta_{\tau-t} \frac{(\tau - t)^2}{(1 + (\tau - (t + 1))\hat{\lambda}_t)^3} \right) - \hat{\lambda}_t \left(\sum_{\tau=t+1}^{\infty} \delta_{\tau-t} \frac{(\tau - t)^2}{(1 + (\tau - (t + 1))\hat{\lambda}_t)^3} \right) \\ & = -\hat{\lambda}_t^2 \left(\sum_{\tau=t+1}^{\infty} \delta_{\tau-t} \frac{(\tau - t)^2}{(1 + (\tau - (t + 1))\hat{\lambda}_t)^3} \right) \leq 0, \end{aligned}$$

where the first inequality follows from the triangle inequality and the second inequality follows from $\delta_{\tau-t} \in [0, 1]$ and $(\tau - (t + 1))\hat{\lambda}_t \geq 0$. Hence, (4) indeed holds for any $\hat{\lambda}_t \in [0, 1]$.

To complete the proof, we first observe that

$$\tilde{f}'(0) = -\frac{\gamma}{2}\nu^2 \sum_{\tau=t+1}^{\infty} \sum_{\ell=t+1}^{\tau} \phi_{\ell-t}\delta_{\tau-\ell} + \nu^2 \frac{\lambda_t}{1-\lambda_t} \Omega = -\frac{\gamma}{2}\nu^2 \sum_{\tau=t+1}^{\infty} \sum_{\ell=t+1}^{\tau} \phi_{\ell-t}\delta_{\tau-\ell} + \nu^2 \hat{\lambda}_{t-1} \Omega,$$

which is strictly negative for any

$$\hat{\lambda}_{t-1} < \frac{\gamma}{2\Omega} \sum_{\tau=t+1}^{\infty} \sum_{\ell=t+1}^{\tau} \phi_{\ell-t}\delta_{\tau-\ell}.$$

By identical arguments as in the proofs of Proposition 3, the agent develops a dogma. Finally, we observe that $\hat{\lambda}_t = 0$ implies $\hat{b}_t = \alpha\gamma\Phi$, which in turn yields $\hat{\mu}_{t+1} - \hat{\mu}_t = \alpha\gamma\Phi$. \square

References

- Brunnermeier, Markus K and Jonathan A Parker**, “Optimal expectations,” *American Economic Review*, 2005, 95 (4), 1092–1118.
- , **Filippos Papakonstantinou, and Jonathan A Parker**, “Optimal time-inconsistent beliefs: Misplanning, procrastination, and commitment,” *Management Science*, 2017, 63 (5), 1318–1340.
- Caplin, Andrew and John V Leahy**, “Wishful Thinking,” https://www.nber.org/system/files/working_papers/w25707/w25707.pdf 2019.