# DIVISIBLE UPDATING

MARTIN W. CRIPPS

Abstract. Two axiomatic characterizations are provided of belief updating. A class of updating processes, termed quasi-Bayesian updating, is characterized by four axioms. These include the divisibility property: that the update does not alter if a composite signal is broken into several parts and several updates. Bayesian updating itself is also characterized using further axioms. Quasi-Bayes updating is applied to a model of sequential sampling. In this model updating that overreacts to new information leads to increased information acquisition.

## 1. Introduction

In this paper we consider the process of updating beliefs in the light of statistical evidence. We treat an updating process as a deterministic map from a prior belief and an information structure (a statistical experiment with a finite number of signals) to a profile of updated beliefs (one for each possible signal in the experiment). We show that an updating process satisfies four axioms if and only if it is a simple generalization of Bayesian updating. Bayesian updating is also characterized: Updating Bayesian if and only if, the above four axioms hold, the updating is continuous, preserves certainty, and the initial belief is required to lie in the convex hull of the updates.

The axiomatic approach to updating taken here is inspired by the axiomatic models of the value of information. In Shannon and Weaver (1949) entropy is derived as the unique measure of information that satisfies certain properties.[1] Our aim is to approach belief revision from a similar normative standpoint, that is, to provide a set of criteria that an agent's belief revision might satisfy if they take "sufficient trouble" to investigate it by some process of introspection.[2] Then, for a given set of criteria, to characterize the class of updating processes that satisfy these conditions. The criteria studied here are not claimed to be definitive, the

---

*Date*: Originally 2019. This version June 19, 2020.

Department of Economics, University College London; m.cripps@ucl.ac.uk. My thanks are due to three referees and the editor for their huge contribution to this paper. Also to Aislinn Bohren, Philippe Jeheil, Alvaro Sandroni, and Ran Speigler for their suggestions and comments. Finally I should thank seminar audiences at Caltech, MIT-Harvard, PSE and Penn.

[1]See Pomatto, Strack, and Tamuz (2018) for a recent contribution to this literature.

[2]This borrows from (Savage, 1972, p.20) and his motivation for the normative use of his postulates.

aim is just to begin to understand the properties of the many models of updating and belief formation that are present in the literature. None of the criteria we consider below are novel. The one we emphasize is a commutative property that we call divisibility. We also study the properties of Bayesian plausibility, which requires that belief updating is an unconditional martingale.[3]

Consider an agent who receives two pieces of information. There are several ways that she can use these two signals to update her beliefs about the world. One is to consider the joint distribution of the two signals and to do just one update. An alternative (which is natural when the signals arrive sequentially) is for her to update her beliefs twice. That is, for her to update her beliefs once using the first signal and then to update her intermediate beliefs using the second signal and its conditional distribution. If these two different procedures generate the same final profile of updated beliefs, we will say that the updating is divisible. Features of this example are made more precise in our definition of divisibility, but the idea at its core is that breaking down composite signals and doing many updates does not affect the eventual updated beliefs.

The main result of this paper shows that updating satisfies four Axioms, if and only if, it is characterized by the process shown in Figure 1, where $F$ is an arbitrary bijection. The four axioms are: First, that uninformative experiments do not result in changes in beliefs. Second, that the names of the signals do not affect the updating just their probability content. Third, that all updated beliefs are possible given the right information structure. And fourth, the divisibility property described above. The updating procedures that satisfy these four properties follow the steps that are illustrated in the figure below. The initial beliefs are mapped to a "Shadow Prior" using a bijection $F$. Then, these shadow priors are updated in the standard Bayesian fashion using the statistical information that is observed to create a Shadow Posterior. Finally, the shadow posterior is mapped back to the space of original beliefs using the inverse bijection $F^{-1}$.

$$\begin{tikzcd} \text{Beliefs} \arrow[d, "F"] & & \text{Updated Beliefs} \\ \text{Shadow Prior} \arrow[rr, "\text{Bayes Updating}"] & & \text{Shadow Posterior} \arrow[u, "F^{-1}"'] \end{tikzcd}$$

FIGURE 1. Quasi-Bayesian Updating, or Updating that Satisfies Axioms 1–4.

Bayesian updating is clearly an element of this class of updating processes ($F$ is the identity). The geometric probability weighting model of updating of Angrisani, Guarino, Jehiel, and Kitagawa (2017) and Bohren and Hauser (2017) is also in this class. Specific properties of the bijection $F$ will determine different properties of the belief revision. We will show that drift in

[3]See Kamenica and Gentzkow (2011) for example.

the updating, overreaction and under-reaction to signals can all be generated by appropriate choices of $F$.

Many of the useful properties of Bayesian updating also carry over to the class of quasi-Bayesian updating processes described by Figure 1. Under Bayesian updating the beliefs are a martingale, they converge, and they are consistent (in correctly specified models). Although the actual beliefs described in Figure 1 are not a martingale, the shadow beliefs are. Thus, if the function $F$ is continuous and maps the extreme points of the belief simplex to themselves, then consistency will hold for this larger class of updating processes. The relative ease with which consistency is established in this class of updating processes contrasts with other models of non-Bayesian updating, for example Rabin and Schrag (1999), Bohren and Hauser (2017) and Epstein, Noor, and Sandroni (2010) where it is much harder to establish.

The simple nature of quasi-Bayesian updating also allows it to be applied to a large class of learning problems. In an extended example, we apply quasi-Bayesian updating to the classic sequential sampling problem of Wald (1945). The value of acquiring a signal is intimately related to the way an agent processes this signal in their updating. Thus the dynamic cost of information will be sensitive to changes in the mode of updating. In this example, agents with updating that is responsive to signals learn more and are willing to buy more signals. Conversely, agents whose updating is unresponsive to signals choose to acquire fewer signals. They do not compensate for their slow learning by acquiring more signals. Thus the responsiveness of the updating and the acquisition of signals are complements.

A natural question to ask is: what normative properties characterize Bayesian updating? We will show the updating is Bayesian if and only if it satisfies: the above four axioms, continuity, $F$ maps the extreme points of the belief simplex to themselves, and one further condition. This condition is that the original belief is contained in the convex hull of the set of updated beliefs. When there are more than two states, we show that these conditions characterize Bayesian updating. Clearly, if the original belief equals the expected value of the updated belief then our condition on the convex hull is also satisfied. Thus, the martingale property, unbiased learning, or Bayes plausibility are all sufficient to ensure that the updating characterized by our four axioms is Bayesian. We also investigate the role that Bayes plausibility plays in discipling the updating when no other axioms are imposed. We show that updating with this property is equivalent to that of a Bayesian who has misspecified probabilities of signals. Thus, one can interpret this kind of updating as the behavior of someone who is a Bayesian but has incorrect knowledge of the process generating the signals.

Why emphasize the divisibility of updating rather than some other property it may have? Our first motivation for the divisibility property is best understood by way of an example. Consider an agent with an arbitrary updating process who independently samples the signals $s \in \{-1, 0, +1\}$. This agent has chosen her updating to have the property that an uninformative signal (one with the same probability in each state) does not change her beliefs, because she thinks this is a sensible property for her updating to have. However, she has not chosen her updating to be divisible. In state $\theta$ the signals, $s$, are sampled with probabilities $(p, \frac{1}{2}, \frac{1}{2} - p)$ and in state $\theta'$ the order of the probabilities are reversed to $(\frac{1}{2} - p, \frac{1}{2}, p)$, where $p < \frac{1}{4}$. Clearly, if she samples $s = 0$ her updated beliefs will not change, because this signal is uninformative. What if she samples two signals and observes the pair $(-1, 1)$? She might process this information iteratively: first updating using the $s = -1$ signal and then updating again using the $s = +1$ signal. The first update could lead her belief in the state $\theta'$ to increase as $\frac{1}{2} - p > p$. And the second update might lead her belief in $\theta'$ to fall back. There is no property of the updating that speaks to the relative size of these two effects, so the agent might end this process with higher, or lower, belief in $\theta$. However, the pair $(-1, 1)$ is uninformative and occurs with probability $p(\frac{1}{2} - p)$ in both states. Moreover, her updating was required to be unaffected by such uninformative signals. Thus this way of processing information violates a property that she intended her updating to have. If she had done one update on this pair, then her beliefs would not have changed. There is an infinite set of uninformative signals she can see—any collection of signals with an equal number of $+1$'s and $-1$'s. If she processes the information in these signals iteratively, then she risks treating such signals as being informative. Indeed if sufficiently long vectors of signals occur, it is possible that an uninformative signal will be treated by this agent as arbitrarily informative. The original property that she wanted her updating to satisfy (uninformative signals result in no change in belief) can be entirely voided if the updating is not divisible. Imposing the additional requirement of divisible updating ensures that all uninformative signals, no matter how they are processed, are recognized to be uninformative. In summary, if an agent believes that their updating ought to satisfy a property such as non-responsiveness to uninformative signals, then divisibility is a method of ensuring the updating consistently obeys this property. Without divisibility it is not clear what any property placed on the updating actually achieves.[4]

In response to the above example, one might argue that the fault is with the agent for choosing to update iteratively rather than in one round. However, there are many reasons updating might proceed iteratively: The signal may arrive in pieces at different points of time. The signal may come from several distinct sources. It may reduce the cognitive load if a signal is processed in several steps because the signal is so complex. Plott (1973) termed this a "divide and conquer strategy".[5] He argued that path independence was a weaker but necessary property of rational social choice. One normative argument he gave in favor of path independence was based on the idea that, "The status quo, or history, should play no dominant role in the determination of choice" (Plott, 1973, p.1087). A similar claim is made here for belief updating. Belief formation ought to be independent of the order that the evidence is processed, because if it was not independent then the status quo (or initial belief)

[4]The example focusses on the uninformative signals, other properties (such as monotonicity) can be voided if divisibility is not required. In Section 5.2 we, further, show that a sophisticated agent can exploit any non-divisibility to maximize what they learn from their signals.

[5]My thanks to a referee for bringing this paper to my attention.

plays a special role in the determination of updated beliefs.[6] To see this, consider an agent who has: the prior belief $\mu$, seen a signal $s$, and has formed the updated beliefs $\mu'$. She might be concerned that her updating was forcing her initial beliefs, $\mu$, to play a special role in the determination of $\mu'$. One way for her to verify that this is not the case would be to suppose that $\mu$ had itself been generated by some previous evidence $s''$ at the previous beliefs $\mu''$. If she circumvented the belief $\mu$ and just saw $s''$ and $s$ at $\mu''$ would she still have the final beliefs $\mu'$? This test of the initial beliefs is obviously identical to the divisibility property. It also ascertains that the initial belief is not being treated in a special way, because it is the initial belief. Importantly, this test *does not* require the agent to weight equally her initial belief and the signal when performing an update. For example, one of the many models of updating that passes the above test is complete dogmatism, that is, for her never to alter her belief whatever evidence she sees. The test described above instead requires that initial beliefs interact in a consistent way with the updating.

The argument presented in the previous paragraph says nothing about preferences. There is considerable evidence from psychology that individuals do have preferences about the order in which signals arrive and are processed, see Legg and Sweeny (2014) for example.[7] These preferences are also studied in the economics literature, Ely, Frankel, and Kamenica (2015) for example. These literatures focus on how the order of signals affects preferences. They do not say that order of signals necessarily effects agents' ultimate beliefs in any way. Divisibility is the claim that the order of signals does not affect their ultimate beliefs.

Another motivation for divisibility comes from outside Economics. In the forward of Ramsey (1926), the theory of probability is described as, "...a branch of logic, the logic of partial belief and inconclusive argument". There is a literature on the logic of partial belief in theories. In particular on how this belief should be revised or updated in the light of new evidence. This literature has considered many forms of belief revision based on the axioms of probability or logic.[8] In most logics, the connective "and" is commutative, that is, exactly the same propositions can be deduced from $A \wedge B$ or $B \wedge A$. Divisibility extends this property to the domain of partial belief. If belief formation was not divisible, then the beliefs after $A$ then $B$ could differ from the beliefs after learning $B$ then $A$. Hence, divisibility under another name has been proposed to a necessary property of belief revision in this literature.[9]

This paper is organized as follows. In Section 2 our lead examples of divisible and non-divisible updating are discussed and there is brief description of the literature. Section 3 provides the formal model of updating and describes the axioms we impose on it. The main result, the characterization of divisible updating, is given in Section 4. Section 5 then returns

[6]The analogy made here between choice with reference points and updating is not new. It has been noted by Rubinstein and Zhou (1999) among others.

[7]There is also genre of humor based around getting good and bad news in various orders.

[8]See Howson and Urbach (2006) and Lewis (1976) and may others.

[9]Perea (2009) applies a version of divisibility to the "imaging" model of updating due to Lewis (1976) .

to our examples of divisible and non-divisible updating and shows how these models can be applied in a sequential sampling model. In Section 6 some basic properties of divisible updating are described. In Section 7 a characterization of Bayesian updating is provided.

## 2. Examples and the Literature

### 2.1. *Examples of Divisible and Non-Divisible Updating*

In this section we consider a simple learning problem and give a brief discussion of our principal examples of divisible and non-divisible updating. A more detailed investigation of these two examples is delayed until Section 5, although some of the results in that section are described here. The example of non-divisible updating is due to Epstein, Noor, and Sandroni (2010). The example of divisible updating is the geometric weighting scheme used by many including Angrisani, Guarino, Jehiel, and Kitagawa (2017).

Consider an agent who is waiting for a bus, but who does not know the arrival process of buses. There are two states for the arrival process: In the *good* state a bus will arrive in period $t = 0, 1, \dots$ with probability $(1-\alpha)\alpha^t$ ($\alpha \in (0,1)$) while in the *bad* state a bus arrives in period $t$ with probability $(1-\beta)\beta^t$ where $\alpha < \beta$. She has initial beliefs $\mu_0 \in (0,1)$ that the state is good. If no bus arrives in period $t = 0$, then a Bayesian would revise her beliefs in the good state downward to $\frac{\alpha\mu_0}{(1-\mu_0)\beta+\mu_0\alpha}$. If no bus arrives after $t$ periods of waiting then her beliefs in the good state (at the start of period $t$) fall to $\frac{\alpha^t\mu_0}{(1-\mu_0)\beta^t+\mu_0\alpha^t}$.

Epstein, Noor, and Sandroni (2010) consider a model of non-Bayesian learning where, if a bus does not arrive, the agent revises her beliefs to a weighted average of her prior and the Bayesian posterior:

$$\mu_1 = (1-\lambda)\mu_0 + \lambda\frac{\alpha\mu_0}{(1-\mu_0)\beta+\mu_0\alpha}, \qquad \lambda \geq 0. \tag{1}$$

This is a particularly elegant model of updating as the linear combination of prior and posterior preserves the martingale property of Bayesian updating when $\lambda < 1$. Choosing $\lambda < 1$ also allows the agent to be conservative in her updating, Hagmann and Loewenstein (2017). Conversely, $\lambda > 1$ allows her to overemphasize the new information she has received.

This updating is non-divisible, so there are many possible values the agent's beliefs could take after $t$ periods waiting for a bus. Iterating (1) for each of the $t$ periods she has been waiting, gives the updated belief

$$\mu_\tau = (1-\lambda)\mu_{\tau-1} + \lambda\frac{\alpha\mu_{\tau-1}}{(1-\mu_{\tau-1})\beta+\mu\alpha}, \qquad \tau = 0, 1, \dots, t.$$

An alternative application of (1) would be to do one update of her prior, $\mu_0$, using all the information accumulated by time $t$ . This would give an updated belief:

$$\hat{\mu}_t := (1-\lambda)\mu_0 + \lambda \frac{\alpha^t \mu_0}{(1-\mu_0)\beta^t + \mu_0 \alpha^t}.$$

In Lemma 3 we will compare $\hat{\mu}_t$ and $\mu_t$. We show that $\hat{\mu}_t < \mu_t$ when $\lambda > 1$ and $\mu_0$ is small. If the agent believes the bad state is likely and overweights new information, then she becomes pessimistic faster if she does one update rather than updating iteratively. This effect is reversed when $\mu_0$ is high, however. A sophisticated agent could achieve the most rapid learning about the state by first updating iteratively when beliefs are high and then doing one further update when her beliefs hit a threshold.

Our main example of divisible updating is a simple generalization of Bayes rule, where the probabilities in the updating are weighted by raising them to a power

$$\mu_\tau = \frac{\alpha^{\frac{1}{a}} \mu_{\tau-1}}{\mu_{\tau-1}\alpha^{\frac{1}{a}} + (1-\mu_{\tau-1})\beta^{\frac{1}{a}}}, \qquad a \geq 0; \quad \tau = 1, \dots, t.$$

We call this geometric probability weighting. This divisible generalization of Bayesian updating appears in Grether (1980), Hanany and Klibanoff (2009), and Angrisani, Guarino, Jehiel, and Kitagawa (2017) for example. In Section 4.1 we show that this model can be interpreted as a geometric average of the prior and Bayesian posterior. Also, that it is the only divisible updating process that reweights probabilities. Other examples of divisible updating are given Sections 4.2, 4.3 below.

### 2.2. *Related Literature*

There is a large and growing literature, both experimental and theoretic, investigating the consequences of a non-Bayesian updating of beliefs, see for example: Rabin and Schrag (1999), Ortoleva (2012), Levy and Razin (2017) among many others. Much of this literature combines issues of updating and decision taking. This is not what the current paper does—it focusses solely on the revision of beliefs and the properties one might want to place on this revision. One theme of this literature has been to investigate the consequences of a particular assumption about the updating. The aim here is somewhat different, that is to try to understand what updating procedures are consistent with a given property. One exception to the focus on decision taking is Epstein, Noor, and Sandroni (2010) who provide a model of updating that captures the under and overreaction to new information. Their model of updating is distinct from the class we consider in several respects and will be considered at length below.

In Gilboa and Schmeidler (1993) the notion of divisibility (termed there commutativity) is introduced and is argued to be an important feature of belief updating particularly in the context of updating ambiguous beliefs. Hanany and Klibanoff (2009), describe a model of updating that has perhaps the closest connection with the one described here. They show that there is a unique "reweighted Bayesian update" that generates a given set of dynamically consistent preferences. They moreover show that this rule satisfies commutativity. These reweighted Bayesian updates are a subset of the class of updating rules that are described in Proposition 1. In Zhao (2016) a set of weaker axioms are shown to characterize an updating rule that does not satisfy divisibility, but does satisfy an order independence property, however, this property is required to hold only for independent events. Similarly, Frankel and Kamenica (2019) consider and order-independence property but here it is required to hold for an expectation rather than for the entire profile of updated beliefs.

There are other characterizations of Bayesian updating in the literature. Majumdar (2004) imposes more regularity (such as monotonicity) and structure on the updating than is required here. Chauvin (2019) considers belief revision in a more abstract setting, without the machinery of statistical experiments. Instead agents observe arguments that map their state to a revised state. Furthermore, neither of the above characterize divisible updating directly.

Several of the updating models described above take the usual formula for Bayesian updating but re-scale the probabilities that appear in it. In contrast, Figure 1 can be viewed as describing updating where the beliefs that appear in Bayesian updating are rescaled, not the probabilities. It is this belief rescaling that distinguishes the generalization of Bayesian updating considered here from much of the literature. We also provide a result, Lemma 2, that describes when these two approaches are equivalent.

## 3. A Model of Belief Updating and the Axioms

In this section we define the updating process $\mathcal{U}$ and the axioms that are imposed on it.

There is an unknown state $\theta$ that can take finitely many values $\theta \in \{1, 2, \ldots, |\Theta|\} := \Theta$. An agent has the full-support beliefs $\mu = (\mu^1, \ldots, \mu^{|\Theta|}) \in \Delta^o(\Theta)$ about this state.[10] The agent also has access to a statistical/Blackwell experiment, $\mathcal{E}$, that provides information about $\theta$. The experiment consists of a finite set of signals $s \in \{1, 2, \ldots, n\} = S_n$ and state-dependent full-support probability distributions for the signals $p^\theta = (p_1^\theta, \ldots, p_n^\theta) \in \Delta^o(S_n)$. The number of signals in the experiment is an arbitrary finite number, thus $\mathcal{E} \in \cup_{n=2}^{\infty} \Delta^o(S_n)^{|\Theta|} := \mathfrak{E}$. We write $\mathcal{E}_n$ to represent an experiment that has $n$ signals.

An updating process, $\mathcal{U}$, takes every belief–experiment pair, $(\mu, \mathcal{E}_n) \in \Delta^o(\Theta) \times \mathfrak{E}$, and maps them to a profile of updated beliefs, $(\mu^1, \ldots, \mu^n) \in \Delta^o(\Theta)^n$. This profile specifies an updated

[10] $\Delta(\Theta)$ denotes the set of probability distributions on the finite set $\Theta$ and $\Delta^o(\Theta)$ denotes the interior of $\Delta(S)$.

belief for each of the $n$ signals in the experiment. Thus $\mathcal{U}$ is described by a sequence of functions $\mathcal{U} := (\mathcal{U}_n)_{n=2}^{\infty}$ where

$$\mathcal{U}_n : \Delta^o(\Theta) \times \Delta^o(S_n)^{|\Theta|} \to \Delta^o(\Theta)^n, \qquad \text{for } n = 2, 3, \dots .$$

We write the elements of this profile as $(\mathcal{U}_n^1, \dots, \mathcal{U}_n^n) = \mathcal{U}_n(\mu, (p^\theta)_{\theta\in\Theta}) = \mathcal{U}_n(\mu, \mathcal{E}_n)$.[11]

We now state some axioms for the updating process $\mathcal{U} := (\mathcal{U}_n)_{n=2}^{\infty}$. The first axiom says that if an experiment is uninformative, then there is no updating. An experiment $\mathcal{E}_n = (p^\theta)_{\theta\in\Theta}$ is uninformative if the signals occur with the same probability in each state, that is, $p^\theta = p^{\theta'}$ for all $\theta, \theta' \in \Theta$. The axiom says that if this is the case, then the elements of the profile of updated beliefs are all equal to the initial beliefs.

**Axiom 1** (Uninformativeness)**.** For all $n$, $\mathcal{U}_n(\mu, \mathcal{E}_n) = (\mu, \dots, \mu)$ if $\mathcal{E}_n = (p^\theta)_{\theta\in\Theta}$ and $p^\theta = p^{\theta'}$ for all $\theta, \theta' \in \Theta$.

The second axiom says that the names of the signals are unimportant for how the beliefs are revised. It is only the probabilities of the signals in the experiment that determine how the beliefs are updated. Thus permuting the order of the signals just permutes the elements of the array of updated probabilities.

**Axiom 2** (Symmetry)**.** For any $n$, any permutation $\omega : \{1, 2, \dots, n\} \to \{1, 2, \dots, n\}$, any $\mu$, and any $\mathcal{E}_n = (p^\theta)_{\theta\in\Theta} \in \mathfrak{E}$,

$$\mathcal{U}_n(\mu, (\omega(p^\theta))_{\theta\in\Theta}) = \left(\mathcal{U}_n^{\omega(1)}(\mu, \mathcal{E}_n), \dots, \mathcal{U}_n^{\omega(n)}(\mu, \mathcal{E}_n)\right),$$

where $\omega(p^\theta) := (p^\theta_{\omega(1)}, \dots, p^\theta_{\omega(n)})$ and $\mathcal{U}_n(.) =: (\mathcal{U}_n^1(.), \dots, \mathcal{U}_n^n(.))$.

At this stage it is useful to introduce a further piece of notation to describe the experiment $\mathcal{E}_n = (p^\theta)_{\theta\in\Theta} \in \mathfrak{E}$. We will define the vector $p_s := (p_s^\theta)_{\theta\in\Theta} \in (0,1)^{|\Theta|}$. This describes the probability of the signal $s$ in each state. The vectors $p^\theta$ are the rows of a Markov matrix while the vectors $p_s$ are its columns. In experiments with only two signals (a binary experiment), the probability of the first signal determines $p^\theta$ completely, so we can write the elements of the array of updated beliefs as follows

$$\mathcal{U}_2(\mu, \mathcal{E}_2) \equiv \left(\mathcal{U}_2^1(\mu, p_1, \mathbb{1} - p_1), \mathcal{U}_2^2(\mu, \mathbb{1} - p_2, p_2)\right), \qquad p_1, p_2 \in (0,1)^{|\Theta|}.$$

If $p_1$ is the vector of state-dependent probabilities for the signal $s = 1$, then $\mathbb{1} - p_1$ are these probabilities for $s = 2$.[12]

[11] In this model the updated beliefs are a deterministic function of the signal and experiment. This is not consistent with all models of updating. For example in Rabin and Schrag (1999) the updated beliefs are randomly determined by a bias that is realized after the signal is observed. To capture this model of updating it would be necessary for the function $U_n$ to take values in $\Delta(\Delta(\Theta))^n$.

[12] We use $\mathbb{1}$ to denote the vector $(1, 1, \dots, 1)$ of appropriate length.

Axiom 3 defines divisibility. It is intended to capture the property that an update is independent of whether information is processed in an iterative way or as one-off process. Consider two possible ways of learning the signal $s \in S_n$. The first is a one-step process where the signal $s$ is generated according to the experiment $\mathcal{E}_n = (p^\theta)_{\theta\in\Theta}$ and then revealed to the agent. Another way of learning $s$ is a two-step process where: First, with probabilities $(p^\theta_1, 1 - p^\theta_1)_{\theta\in\Theta} = (p_1, \mathbb{1} - p_1)$ the signal $s = 1$ or the signal $s \neq 1$ is revealed to the agent in a binary experiment. Then, in the case where the outcome $s \neq 1$ was obtained in the first experiment, a signal from the set $\{2, \dots, n\}$ is generated from a second experiment with the objective conditional probabilities $(p^\theta_{-1}(1 - p^\theta_1)^{-1})_{\theta\in\Theta}$. (As usual, $p^\theta_{-s}$ is the vector $p^\theta$ with the $s^{\text{th}}$ element omitted). In what follows we will use $\mathcal{E}_{n-1} := (p^\theta_{-1}(1 - p^\theta_1)^{-1})_{\theta\in\Theta}$ to denote the experiment that occurs conditional on $s \neq 1$.[13]

Axiom 3 says that these two different processes for observing the signal $s$ have no effect on the agent's ultimate profile of beliefs for all the signals $s \in S_n$. This assertion has two distinct claims: First it says that learning the signal is $s = 1$ when there are $n - 1$ other signals has the same effect on the updated beliefs as learning the signal is $s = 1$ when there is a binary experiment. The relative probabilities of the signals that were *not* observed $(s = 2, \dots, n)$ have no role in determining how beliefs will be updated when $s = 1$ is observed. This is part (a) of the Axiom.

The second part of Axiom 3 says that the agent's updated beliefs when they see the signal $s' \neq 1$ in an experiment, that is $\mathcal{U}^{s'}_n(\mu, \mathcal{E}_n)$, are the same as the updated beliefs they have at the end of the two-step process. In this two-step process they first learn the signal was not $s = 1$ and update their beliefs to $\mathcal{U}^2_2(\mu, p_1, \mathbb{1} - p_1)$. Then they learn that the signal was $s'$ from the experiment $\mathcal{E}_{n-1}$ and engage in a further update to $\mathcal{U}^{s'-1}_{n-1}\left(\mathcal{U}^2_2(\mu, p_1, \mathbb{1} - p_1), \mathcal{E}_{n-1}\right)$.[14] This is what part (b) of the Axiom says.

**Axiom 3** (Divisibility)**.** For all $n \geq 3$, $\mathcal{E}_n \in \mathfrak{E}$, and $\mu \in \Delta^o(\Theta)$:

(a) $\mathcal{U}^1_n(\mu, \mathcal{E}_n) \equiv \mathcal{U}^1_2(\mu, p_1, \mathbb{1}-p_1)$; (b) $\mathcal{U}^{s'}_n(\mu, \mathcal{E}_n) \equiv \mathcal{U}^{s'-1}_{n-1}\big(\mathcal{U}^2_2(\mu, p_1, \mathbb{1} - p_1), \mathcal{E}_{n-1}\big)$, $\forall s' > 1$.

Where $\mathcal{E}_{n-1} := (p^\theta_{-1}(1 - p^\theta_1)^{-1})_{\theta\in\Theta}$.

Iterating the two-step process used in the axiom allows us to consider learning the signal over $n - 1$ rounds of updating. (In the $k^{\text{th}}$ round the agent observes an experiment with an outcome $s = k$ or $s > k$.) Divisibility asserts that this iterated process for learning the signal has no effect on the ultimate beliefs. The symmetry property also implies that the order that the signals are revealed has no effect on the eventual profile of updated beliefs. Hence, these axioms imply that the updating of beliefs is independent of the order that information arrives. That is, reversing the order in which two pieces of information arrive has no effect on the

[13] Recall that $p^\theta \in \Delta^o(\Theta)$ and so $1 > p^\theta_1$.

[14] The notation here requires an explanation. The updated beliefs after the signal $s'$ is the $s'-1^{\text{th}}$ component of the profile $\mathcal{U}_{n-1}$ when the first signal is not present. Hence the change in the superscript on $\mathcal{U}_{n-1}$.

ultimate profile of beliefs. The axiom is also implied by an order-reversal property. The two ways of learning the signal $s$ in the axiom amount to either: First observing the true value of $s$ and then observing (the uninformative experiment) that indicates whether $s \neq 1$. Or, to first observing whether $s \neq 1$ and then observing the true value of $s$. Hence when there is symmetry, Axiom 3 is equivalent to an assertion that the updated beliefs are independent of the order that information arrives.[15]

The final axiom considers the updating in a binary experiment. It says that for at least one initial belief $\mu^o$ all updated beliefs are possible after the signal $s = 1$. Furthermore, there is a unique binary experiment (up to homogeneity) that achieves this updated belief.[16]

**Axiom 4** (Non-Dogmatic)**.** There exists $\mu^o \in \Delta^o(\Theta)$ such that for every $\mu \in \Delta^o(\Theta)$ the equation $\mathcal{U}_1^2(\mu^o, p_1, \mathbb{1} - p_1) = \mu$ has a unique solution $p_1 \in \Delta^o(\Theta)$.

Updating satisfying this axiom is "non-dogmatic" as it describes an agent whose belief can take any value, provided they see the right evidence. This axiom would fail if there is a set $\mu$'s for which this equation cannot be solved by any binary experiment. Such a failure has significant consequences when the updating satisfies divisibility. Divisibility implies that the updated belief after one complex (but finite) history of experiments and signals can be collapsed to a one-step binary experiment where the signal $s = 1$ or $s \neq 1$ is revealed. Thus, if the updating satisfies divisibility, there can be no finite sequence of experiments with finite signals where the ultimate belief is in this set—it is a taboo set of beliefs. An example of divisible updating that violates of this axiom, would be completely dogmatic updating where no evidence ever leads to a change in beliefs, $\mathcal{U}_1^2(\mu^o, p_1, \mathbb{1} - p_1) := \mu$ for all $(\mu, p_s)$. The taboo set here is very large.[17]

The axiom, also, requires a degree of sensitivity to evidence as different evidence cannot generate the same posterior. This sensitivity to evidence can be extreme as there is no requirement that the updating is continuous. The axiom would fail if there were two different binary experiments $p_1, p_1' \in \Delta^o(\Theta)$ that generated the same updated beliefs at $\mu_0$. We will show (see footnote 19 below) that if this were the case, and the updating is divisible, then there are many experiments that are treated as uninformative. Thus if divisible updating violates this axiom, there is a large scale lack of sensitivity to certain kinds of information.[18]

[15]This equivalence argument is formalized in the online appendix.

[16]Usually $p_1 \in (0,1)^{|\Theta|}$, but $\mathcal{U}_1^2$ will be shown to be homogeneous degree zero, so we require a unique solution $p_1 \in \Delta^o(\Theta) \subset (0,1)^{|\Theta|}$.

[17]A less extreme example of quasi-Bayesian updating where some beliefs cannot be reached was suggested by a referee. There are: two states $\theta$ and $\theta'$, initial beliefs $(\mu, 1-\mu)$, and signal probabilities $p_s^\theta, p_s^{\theta'} \in [0,1]$. Define the updated beliefs to be $\mu' = \frac{\mu p_s^\theta}{2\mu p_s^\theta + (1-2\mu) p_s^{\theta'}}$ if $\mu \leq \frac{1}{2}$ and $\mu' = \frac{(2\mu-1)p_s^\theta + (1-\mu)p_s^{\theta'}}{(2\mu-1)p_s^\theta + 2(1-\mu)p_s^{\theta'}}$ if $\mu > \frac{1}{2}$. Beliefs never leave the interval $[0, \frac{1}{2}]$ if they start in this interval, no matter how much information is observed.

[18]In the online appendix we examine divisible updating models where this property fails. In these models the dimension of the space of updated beliefs can be smaller than the dimension of the belief space.

This axiom also excludes models of updating where there are fixed costs of contemplation or of belief revision. In such models (Ortoleva (2012), for example) there are sets of experiments for which it is simply not worth revising beliefs, hence there would be many experiments with the update equal to the prior.

## 4. Characterization of Quasi-Bayesian Updating

In this section a proposition is proved that gives a characterization of updating procedures $\mathcal{U}$ that satisfy Axioms 1–4. We will show that any such updating is characterized by a (not necessarily continuous) bijection $F$ that maps beliefs to a shadow prior. Then, this shadow prior is updated by Bayes rule to create a shadow posterior. Finally the shadow posterior is mapped back by $F^{-1}$ to form the agent's updated beliefs. As was illustrated in Figure 1.

The updating in a binary experiment plays an important role in Proposition 1, so we will define a function $u : \Delta^o(\Theta) \times (0,1)^{|\Theta|} \to \Delta(\Theta)$ to represent it:

$$\mathcal{U}_2(\mu, \mathcal{E}_2) \equiv \left(\mathcal{U}_2^1(\mu, p_1, \mathbb{1} - p_1), \mathcal{U}_2^2(\mu, \mathbb{1} - p_2, p_2)\right) \equiv (u(\mu, p_1), u(\mu, p_2)) \,. \tag{2}$$

The function $u(\mu, p_1)$ gives the updated value of beliefs in a binary experiment where $s = 1$ occurs with the probabilities $p_1 \in (0,1)^{|\Theta|}$. Symmetry allows us to write the profile of updated beliefs for binary experiments, (2), in terms of the one function $u$. Lemma 1 shows that the entries of the updating function for any number of signals $s$ are all made up of this function evaluated at the different signal probabilities $p_s$. It also shows that $u$ satisfies a functional equation and is homogeneous degree zero in $p_s$.

Before stating Lemma 1 we introduce some notation that is used throughout this paper. If $x, y$ are vectors of $m$ strictly positive elements, then $x \circ y$ is used to denote Hadamard (element wise) product of the two vectors and $x \circ y^{-1}$ denotes Hadamard division:

$$x \circ y := (x_1 y_1, \ldots, x_m y_m), \qquad x \circ y^{-1} := (x_1/y_1, \ldots x_m/y_m).$$

Using this notation, we can write the Bayesian updated belief, $\mu^B$, given the priors $\mu$ and the vector of signal probabilities $p_s$ as

$$\mu^B = \frac{\mu \circ p_s}{\mu^T p_s}. \qquad \text{(Bayes Rule)}$$

This expression, suitably adjusted, appears in Propositions 1, 5, and 6 below.

**Lemma 1.** *Suppose the updating $\mathcal{U}$ satisfies Axioms 1, 2, and 3, then:*
*(i)* $\mathcal{U}_n(\mu, \mathcal{E}_n) \equiv (u(\mu, p_1), \ldots, u(\mu, p_n))$ *, for all $n$.*
*(ii)* $u(\mu, p_s)$ *is homogeneous degree zero in $p_s$.*
*(iii)* $u(\mu, p_s) \equiv u\left(u\left(\mu, x\right), p_s \circ x^{-1}\right)$ *for all $x \in \mathbb{R}^n$ with $p_s \leq x < \mathbb{1}$.*

All proofs are in the Appendix. An intuition for the Lemma is as follows. For part (i), observe that Axiom 3(a) requires that the update after signal $s = 1$ is independent of the probabilities of the other signals. Symmetry implies that this property must hold for all signals, not just the signal $s = 1$. Part (ii) of the Lemma says that the updated belief after the signal $s$, is homogeneous degree zero in $p_s$, the vector of probabilities for the signals $s$. Suppose that there are three possible signals $S = \{1, 2, 3\}$ and that one signal, say $s = 1$, is equally likely under all states. Recall that Axiom 3(b) considers two cases: One where $s$ is determined in a one-off experiment. The second where $s$ is determined in a two-step process— first an experiment with binary outcomes $s = 1$ and $s \neq 1$ and then $s \in \{2, 3\}$ is determined. Here observing $s = 1$ is uninformative, so by Axiom 1 the first stage of the two-step process leads to no updating of the priors. But in the second-stage experiment, there is increased the relative probabilities of the signals $\{2, 3\}$. The updates after these signals are equal to the updates after the one-step experiment (with low relative probabilities for $\{2, 3\}$). Thus the scaling up of the probabilities had no effect on the updating of beliefs and the updating is homogeneous. The final part of the Lemma (iii) rewrites the divisibility condition in terms of the function $u$.

In Proposition 1 we show that we know a great deal about the form of $u(.)$, because it satisfies:

$$u(\mu, p_s) \equiv u\left(u\left(\mu, x\right), p_s \circ x^{-1}\right), \qquad \forall\ p_s \leq x < \mathbb{1}; \tag{3}$$

from Lemma 1(iii). This functional equation captures the fact that final beliefs are independent of the order information arrives. To find all updating rules that satisfy Axioms 1–4, it suffices to find all functions $u(.)$ that satisfy (3). In solving functional equations, it is usual to impose some further mathematical regularity on the class of functions one is willing to entertain as solutions. This is the role of Axiom 4. As we have explained above, it ensures the solutions we find are global, so the updating does not have taboo sets of beliefs. It also ensures the full dimensionality of the set of updates.[19] To solve functional equations it is not necessary to assume differentiability or continuity. The proposition below transforms (3) into a well-known functional equation that is called the translation equation.[20] This was solved in its multidimensional form by, Aczél and Hosszú (1956). We use their solution to find (4) the set of functions $u$ that satisfy (3). If (4) is compared with the formula for Bayes updating above, it is clear that Bayes updating is being performed on the beliefs $F(\mu)$ with the signal probabilities $p_s$. The outcome of this is then mapped back by $F^{-1}$.

[19]Axiom 4 requires that $u(\mu, .) = \hat{\mu}$ has a unique solution in $\Delta^o(\Theta)$. If there were two solutions $u(\mu, d) = \hat{\mu}$ and $u(\mu, \hat{d}) = \hat{\mu}$, then (3) would imply that $u(\hat{\mu}, d \circ \hat{d}^{-1}) = u(u(\mu, \hat{d}), d \circ \hat{d}^{-1}) = u(\mu, d) = \hat{\mu}$ . Iterating this $u(\hat{\mu}, (d \circ \hat{d}^{-1})^k) = \hat{\mu}$ for all integers $k$. Thus, the experiment with probabilities $(d \circ \hat{d}^{-1})^k$ is uninformative at $\hat{\mu}$ for all $k$. This can lead to a manifold of experiments that are all uninformative and the dimension of the set of updated beliefs being less than the original space of beliefs. Updating with this property is investigated in the online appendix.

[20]A previous application of the translation equation to economics can be found in Sokolov (2011).

**Proposition 1.** The updating $\mathcal{U}$ satisfies the Axioms 1–4, if and only if there exists a bijection $F : \Delta^o(\Theta) \to \Delta^o(\Theta)$ such that

(4)
$$\mathcal{U}_n(\mu, \mathcal{E}_n) \equiv (u(\mu, p_1), \ldots, u(\mu, p_n)) , \qquad \text{for all } n, \text{ where} \qquad u(\mu, p_s) \equiv F^{-1}\!\left(\frac{F(\mu) \circ p_s}{F(\mu)^T p_s}\right) .$$

We now give a brief sketch of the proof of this result. The transformations we impose on (3) reduce the dimension of the variables $(\mu, x, p_s)$, and $u$, by dividing by the last entry. Then we take logarithms of these ratios. This writes the variables and function in terms of log-likelihood ratios as $(\phi, z, z + y)$ and $\tilde{v}$. With this (3) becomes

$$\tilde{v}(\phi, z + y) \equiv \tilde{v}\left(\tilde{v}(\phi, z), y\right) , \qquad \forall \phi, y, z \in \mathbb{R}^{|\Theta|-1}.$$

This is the translation equation. One simple solution to this functional equation is to add the arguments together: that is $\tilde{v}(\phi, x) = \phi + x$. Transforming this equation back from log-likelihoods into probabilities gives Bayesian updating. In general this functional equation tells us about the contours of $\tilde{v}(.)$. This is because if $y + z$ are constant but $z$ and $y$ vary, then the left of this equation stays constant but the arguments of $\tilde{v}$ on the right vary. It is then relatively simple to see that each such contour of the function $\tilde{v}(.)$ is a translation of the other. Thus once the form of one contour has been determined all other contours are just translations of it. Choosing one arbitrary function to determine the shape of a contour and a second to determine the value taken by each contour is sufficient to determine $v$. It will be convenient to have a name for the updating characterized by Proposition 1. We will call it quasi-Bayesian updating.

**Definition 1.** *The updating $\mathcal{U}$ is said to be quasi-Bayesian if it satisfies Axioms 1–4.*

Proposition 1 permits $\mathcal{U}$ to be discontinuous. If $\mathcal{U}$ is continuous on $\Delta(\Theta) \times \Delta(S)^{|\Theta|}$ (the interior and the boundary), then we can restrict the bijection to be a homeomorphism and extend its domain to $\Delta(\Theta)$. This is summarized in the following corollary.

**Corollary 1.** *The updating $\mathcal{U}$ satisfies Axioms 1–4 and is continuous then there exists a homeomorphism $F : \Delta(\Theta) \to \Delta(\Theta)$ such that*

$$\mathcal{U}_n(\mu, \mathcal{E}_n) \equiv (u(\mu, p_1), \ldots, u(\mu, p_n)) , \qquad \textit{for all } n, \textit{ where} \qquad u(\mu, p_s) \equiv F^{-1}\!\left(\frac{F(\mu) \circ p_s}{F(\mu)^T p_s}\right) .$$

We now give four examples of updating rules in the class characterized by Proposition 1. These show how one important model of non-Bayesian updating falls in this class and describe other discontinuous and non-monotonic models of updating.

### *4.1. Geometric Probability Weighting*

In this example the quasi-Bayesian updating is generated by the homeomorphism

$$F^a(\mu) := \left( \frac{\mu_1^a}{\sum_\theta \mu_\theta^a}, \ldots, \frac{\mu_{|\Theta|}^a}{\sum_\theta \mu_\theta^a} \right); \qquad a > 0.$$

We first show how an explicit formula for the updating can be derived from this. First, we rewrite the condition (4) as $F^a(u(\mu, p_s)) = (F^a(\mu) \circ p_s)/F^a(\mu)^T p_s$. This implies

$$\frac{F_\theta^a(u(\mu, p_s))}{F_{\theta'}^a(u(\mu, p_s))} = \frac{p_s^\theta}{p_s^{\theta'}} \frac{F_\theta^a(\mu)}{F_{\theta'}^a(\mu)}. \tag{5}$$

This is particularly useful as it allows us to derive a property for ratios of the updated beliefs without requiring an explicit expression for inverse of the homeomorphism $F^a$.

$$\frac{u_\theta(\mu, p_s)}{u_{\theta'}(\mu, p_s)} = \left( \frac{p_s^\theta}{p_s^{\theta'}} \right)^{1/a} \frac{\mu_\theta}{\mu_{\theta'}}, \qquad \text{and so} \qquad u(\mu, p_s) \equiv \frac{\mu \circ (p_s)^{1/a}}{\mu^T (p_s)^{1/a}};$$

where $(p_s)^{1/a} := \left( (p_s^1)^{1/a}, \ldots, (p_s^{|\Theta|})^{1/a} \right)$. This explicit form for an updating rule is well known and used in Angrisani, Guarino, Jehiel, and Kitagawa (2017) and Bohren and Hauser (2017), for example. Although it appears that the characterization of quasi-Bayesian updating (4) worked by transforming the priors that enter the Bayesian formula, in this case the updating reweights the probabilities in the Bayesian formula not the priors.

The updating generated by $F^a$ can also be interpreted as a geometric weighted-average of the prior and the Bayes update, that is,

$$\frac{\mu \circ (p_s)^{1/a}}{\mu^T (p_s)^{1/a}} = K \underbrace{(\mu)^{1-\frac{1}{a}}}_{\text{prior}} \circ \underbrace{\left( \frac{\mu \circ p_s}{\mu^T p_s} \right)^{\frac{1}{a}}}_{\text{Bayes Rule}}.$$

$K$ is a normalizing constant chosen to ensure the RHS is in $\Delta(\Theta)$. Hence if $a < 1$, the agent overreacts to new information and over weights the Bayesian update. Conversely if $a > 1$, therefore, the agent under-reacts to new information—they place too much weight on their prior and do not adjust their beliefs as much as a Bayesian would.

Geometric probability weighting is a member of the class of updating rules where the probabilities that enter into Bayes rule are rescaled in some way:

$$u^j(\mu, p_s) := \frac{\mu \circ J(p_s)}{\mu^T J(p_s)}, \qquad J(p_s) := \left( j_1(p_s^1), \ldots, j_{|\Theta|}(p_s^{|\Theta|}) \right).$$

(Where: $j_\theta(0) = 0$ and $j_\theta(.)$ is increasing, for all $\theta$.) Geometric probability weighting is the only updating rule consistent with Proposition 1 that is of this form.

**Lemma 2.** $u^j(\mu, p_s) \equiv F^{-1}\left(\frac{F(\mu) \circ p_s}{F(\mu)^T p_s}\right)$ *for some homeomorphism $F$ iff* $u^j(\mu, p_s) \equiv \frac{\mu \circ (p_s)^a}{\mu^T (p_s)^a}$.

The class of updating rules in Proposition 1 generalize Bayes by rescaling beliefs. Probability weighting rules generalize Bayes by rescaling probabilities. Geometric probability weighting is the class of updating rules where these two generalizations of Bayes intersect.

### 4.2. *Exponential Weighting*

The second homeomorphism we consider, $F^b$, generates a divisible updating process that is similar to multinomial logit.

$$F^b(\mu) := \left( \frac{e^{-\frac{b}{\mu_1}}}{\sum_\theta e^{-\frac{b}{\mu_\theta}}}, \dots, \frac{e^{-\frac{b}{\mu_{|\Theta|}}}}{\sum_\theta e^{-\frac{b}{\mu_\theta}}} \right)$$

The calculation of (5), the ratio of the updated beliefs, for this homeomorphism gives

$$\frac{1}{u_{\theta'}(\mu, p_s)} - \frac{1}{u_\theta(\mu, p_s)} = \frac{1}{\mu_{\theta'}} - \frac{1}{\mu_\theta} + \frac{1}{b} \ln \frac{p_s^\theta}{p_s^{\theta'}}.$$

This updating gives a linear shift in the inverse probabilities of each state that is determined by the log likelihood of the signal probabilities. The constant $b$ determines how responsive the update is to the information in the signal. For example, when there are only two states with initial probabilities $(\mu, 1-\mu)$ the formula for the updated beliefs $(\hat{\mu}, 1-\hat{\mu})$ is

$$\frac{2\hat{\mu} - 1}{\hat{\mu}(1-\hat{\mu})} = \frac{2\mu - 1}{\mu(1-\mu)} + \frac{1}{b} \ln \frac{p_s^\theta}{p_s^{\theta'}}.$$

The function $\frac{2\mu-1}{\mu(1-\mu)}$ is increasing so when $p_s^\theta > p_s^{\theta'}$ updated beliefs move upwards. For large values of $b$ this updating exhibits under-reaction to signals, that is, it overweights the prior (see Section 6.3).

### 4.3. *Two Counter Examples*

Now two quasi-Bayesian updating rules are described for $\Theta = \{\theta, \theta'\}$ that have strange or undesirable properties. The first we call bad Bayesian updating. This is updating that concludes the state is $\theta$ when there is a signal that only happens in state $\theta'$. Consider the bijection that exchanges the probabilities of the states $F : (\mu_\theta, \mu_{\theta'}) \mapsto (\mu_{\theta'}, \mu_\theta)$. In this case the updated beliefs are

$$u_\theta(\mu, p_s) = \frac{\mu_\theta p_s^{\theta'}}{\mu_\theta p_s^{\theta'} + \mu_{\theta'} p_s^\theta}.$$

This updating applies Bayes rule incorrectly—it uses the signal probabilities associated with the state $\theta'$ to weight the state $\theta$. An agent who updates in this way will increase their belief in the state $\theta$ when they see a signal that is more likely in state $\theta'$. In particular, if $p_s^{\theta'} > 0$ and $p_s^\theta \to 0$, then the update above gives probability one to the state $\theta$ although the signal $s$ only arises in state $\theta'$.

A second oddity is generated by the discontinuous non-monotonic bijection

$$F^d(\mu_\theta, \mu_{\theta'}) := \begin{cases} (\frac{1}{2} - \mu_\theta, \frac{3}{2} - \mu_{\theta'}) & \text{if } \mu_\theta \in (0, \frac{1}{2}), \\ (\frac{1}{2}, \frac{1}{2}) & \text{if } \mu_\theta = \frac{1}{2}, \\ (\frac{3}{2} - \mu_\theta, \frac{1}{2} - \mu_{\theta'}) & \text{if } \mu_\theta \in (\frac{1}{2}, 1). \end{cases}$$

$F^d$ is its own inverse, so an explicit expression for the updated beliefs exists. The updated belief is given below in the case where $p_s^\theta < p_s^{\theta'}$ and $\mu_\theta < \frac{1}{2}$.

$$u_\theta(\mu, p_s) = \frac{1}{2} - \frac{(\frac{1}{2} - \mu_\theta)p_s^\theta}{(\frac{1}{2} - \mu_\theta)p_s^\theta + (\frac{3}{2} - \mu_{\theta'})p_s^{\theta'}} > \mu_\theta.$$

This example will be useful when we discuss the role of continuity in proving the consistency of divisible updating.

## 5. Divisible Updating and The Wald Problem

In this section it is shown that the class of quasi-Bayesian updating rules in Proposition 1 can be readily applied to a classic model of sampling (Wald (1945)). The Wald sequential sampling model is the most important dynamic model of costly information acquisition.[21] It is obvious that updating affects agents' information acquisition decisions, here we show that quasi-Bayesian updating is a natural way of studying this relationship. We show, in an example, that the amount of sampling increases as the updating used becomes more responsive to current signals. Agents who overweight their signals tend to collect more information and, as a result, become better informed. They, also, value information more highly and pay higher sampling costs. Thus information acquisition decisions have a simple relationship with the properties of the updating. Sequential sampling with non-divisible updating is also considered. There is no simple application of this model of updating to sequential sampling. For example, we show how the agent can, then, benefit from committing to ignoring their signals for a predetermined time interval and describe their sampling behavior in this case.

The model of signals in this section is different from the one above, as we use a continuous-time model of sampling. It is a continuous-time version of the lead example in DeGroot (1970), Chapter 12.[22] There are two states $\theta \in \{1, -1\}$. At time $t \geq 0$ the agent can pay a flow cost $c\,dt$ to observe a state-dependent signal process for a time interval $dt$: In state $\theta = 1$, a Poisson process generates the signal $s_t = 1$ with the arrival rate $\alpha\,dt$. In state $\theta = -1$, a Poisson process generates the signal $s_t = -1$ with the arrival rate $\beta\,dt < \alpha\,dt$. The arrival of a non-zero signal reveals the state. If there is no arrival we define the signal to be $s = 0$

[21] It is used, by Fudenberg, Strack, and Strzalecki (2018) Morris and Strack (2017), for example, to model how agents form beliefs and act.

[22] More complex models of Poisson sampling are studied by Nikandrova and Pancs (2018) and Che and Mierendorff (2019), for example.

over the interval $dt$. This zero signal is also informative, because the arrival rate of signals is highest in state $\theta = 1$.[23]

The agent's decision problem is to decide how much data to sample before irrevocably choosing an action $\xi \in \{-1, 1\}$. They incur flow costs from their sampling of $c\,dt$ and they incur a lump-sum loss normalized to 1 if their irrevocable action does not match the state (and a loss of 0 otherwise). If they have beliefs $\mu$ that $\theta = 1$, then their expected loss from the optimal immediate action is $\min\{\mu, 1-\mu\}$.

### 5.1. *Divisible Updating*

We consider an agent with quasi-Bayesian updating who faces this sequential sampling problem. First, we show that their value function for optimal sampling is a simple transformation of that of a Bayesian. Then, a complete solution to the optimal sampling problem for the geometric weighting version of quasi-Bayesian updating is given. In this solution we find that responsiveness to information and signals are complementary, that is, as the agent's updating becomes more responsive to their current signal the amount of sampling they choose to do increases.

Let us begin by describing Bayesian updating. At $t = 0$ the agent has the belief $\mu_0$ that $\theta = 1$. If she sampled signals over the time interval $[0, t]$, then with probability $e^{-\alpha t}$ (respectively $e^{-\beta t}$) there is no arrival in state $\theta = 1$ (respectively $\theta = -1$) and she would form the Bayesian update

$$\mu_t = \begin{cases} 1 & \text{if } s_\tau = 1 \text{ for some } \tau \leq t; \\ \frac{\mu_0 e^{-\alpha t}}{\mu_0 e^{-\alpha t} + (1-\mu_0)e^{-\beta t}} & \text{if } s_\tau = 0 \text{ for all } \tau \leq t; \\ 0 & \text{if } s_\tau = -1 \text{ for some, } \tau \leq t. \end{cases}$$

It is simple to describe the beliefs of an agent with quasi-Bayesian updating. Let $f : [0,1] \to [0,1]$ be the increasing homeomorphism that maps $(\mu, 1-\mu)$ to her "shadow beliefs" $(f(\mu), 1-f(\mu))$. Define $\nu_t := f(\mu_t)$ to be the value of these shadow beliefs. By Proposition 1, $\nu_t$ is updated using Bayes rule:

$$(6) \qquad \mu_t = f^{-1}(\nu_t), \text{ where } \nu_t = \begin{cases} 1 & \text{if } s_\tau = 1 \text{ for some } \tau \leq t; \\ \frac{f(\mu_0) e^{-\alpha t}}{f(\mu_0) e^{-\alpha t} + (1-f(\mu_0))e^{-\beta t}} & \text{if } s_\tau = 0 \text{ for all } \tau \leq t; \\ 0 & \text{if } s_\tau = -1 \text{ for some, } \tau \leq t. \end{cases}$$

Actual beliefs satisfy $\mu_t \equiv g(\nu_t)$ where $g := f^{-1}$, so we will treat $\nu_t$ as the state variable in the sequential sampling problem. For example, $g(\nu_t)$ is the probability $\theta = 1$ at the state $\nu_t$ and the expected loss from immediate action is $\min\{g(\nu_t), 1 - g(\nu_t)\}$.

[23] This is also a continuous time version of the model of the arrival of buses in Section 2.1. The only difference is that here when a bus arrives the agent learns the state for certain.

The agent's beliefs $g(\nu_t)$ decrease if she only observes the zero signal in her sampling ($f(.)$ is assumed to be increasing). Hence, she must decide how long she should continue with her sampling before acting. Her decision problem is, therefore, described by a stopping time $\tau \geq 0$. If she prefers to act immediately without sampling, then $\tau = 0$. Or if $\tau > 0$, she prefers to sample for $\tau$ periods with no non-zero signal before choosing an optimal action.

We now describe the agent's optimal sampling problem, her value function, and her HJB equation. At time $t$ the agent believes a non-zero signal arrives with probability $\alpha g(\nu_t)dt + \beta(1-g(\nu_t))dt$. Hence, she attaches probability $P_t := g(\nu_0)e^{-\int_0^t \alpha g(\nu_s)ds} + (1-g(\nu_0))e^{-\int_0^t \beta(1-g(\nu_s))ds}$ to the event that there are $t$ periods of sampling with only the zero signal. Her expected payoff from following an optimal sampling policy at state $\nu_0$ is

$$\mathcal{L}(\nu_0) := \min_{\tau \geq 0} \int_0^\tau cP_t \, dt + P_\tau \min\{g(\nu_\tau), 1 - g(\nu_\tau)\}.$$

As $\tau$ is a deterministic threshold, this minimization problem can be solved by using simple calculus. However, treating the shadow beliefs as the state variable, we can also write the HJB equation for this minimization

$$\mathcal{L}(\nu) = \min\{\, \min\{g(\nu), 1 - g(\nu)\}\, ,\, c\,dt + [1 - \alpha g(\nu)dt - \beta(1 - g(\nu))dt]\, \mathcal{L}(\nu + d\nu)\,\}.$$

The second term in the minimum sums the agent's expected loss from sampling $dt$ more periods and her continuation value. When a signal arrives (with probability $\alpha g(\nu)dt + \beta(1 - g(\nu))dt$) the state is revealed, losses are zero, and no more sampling is required. If not, the shadow beliefs are updated to $\nu + d\nu$ and there is a new continuation value $\mathcal{L}(\nu + d\nu)$. The usual formula for the Bayesian updating of Poisson processes applies to $\nu$, that is, $d\nu = (\beta - \alpha)\nu(1 - \nu)\, dt$. If it is optimal to continue sampling (the minimum is attained by the second expression in the braces), the loss function therefore, satisfies the ODE

$$\mathcal{L}'(\nu) + \frac{g(\nu) + y}{\nu(1-\nu)}\mathcal{L}(\nu) = \frac{x}{\nu(1-\nu)}, \qquad y := \frac{\beta}{\alpha - \beta}, x := \frac{c}{\alpha - \beta}. \tag{7}$$

When $g$ is the identity, the loss function $\mathcal{L}(\nu)$ describes a standard optimal sampling problem with Bayesian updating. If $2c < \beta$, the solution to the optimal sampling problem with Bayesian updating is described by two thresholds $\underline{\mu} < 1/2 < \bar{\mu}$. For $\mu \geq \bar{\mu}$ the immediate action $\xi = +1$ is optimal and for $\mu \leq \underline{\mu}$ the action $\xi = -1$ is optimal. For $\mu$ in the interval $(\underline{\mu}, \bar{\mu})$ the agent samples and, if no revealing signal arrives, the beliefs drift down.

These properties of the Bayesian solution extend to settings where the agent is not Bayesian but does update divisibly. Proposition 2 describes the optimal policy and value function of an agent who uses geometric probability weighting: $\mu = g(\nu) = \frac{\nu^{1/a}}{\nu^{1/a}+(1-\nu)^{1/a}}$. The proposition also gives a solution for the cutoff $\underline{\mu}_a$, where the sampling stops. The value function and cutoffs are illustrated in Figure 2 where: the blue lines describe the losses from immediate action, the red lines the expected losses from optimal sampling, and the cutoffs $(\underline{\mu}_a, \bar{\mu}_a) = (g(\underline{\nu}_a), g(\bar{\nu}_a))$ are where they intersect.

**Proposition 2.** If: $\alpha > \beta > 2c$, the agent has quasi-Bayesian updating with $\mu = g(\nu) = \frac{\nu^{1/a}}{\nu^{1/a}+(1-\nu)^{1/a}}$, and $a > 0$. Then, there exists $\underline{\nu}_a \in (0, \frac{1}{2})$ and $\bar{\nu}_a \in (\frac{1}{2}, 1)$ such that the optimal sampling policy to sample iff $g^{-1}(\mu) = \nu \in (\underline{\nu}_a, \bar{\nu}_a)$. For $\nu \in (\underline{\nu}_a, \bar{\nu}_a)$:

$$\mathcal{L}(\nu) = \frac{x \int_{\underline{r}_a}^{r} \rho^{y-1}(1+\rho^{1/a})^a \, d\rho + K_a}{r^y(1+r^{1/a})^a}, \qquad r := \frac{\nu}{1-\nu}, \underline{r}_a := \frac{\underline{\nu}_a}{1-\underline{\nu}_a}. \tag{8}$$

And: $\mathcal{L}(\nu) = g(\nu)$ for $\nu < \underline{\nu}_a$, $\mathcal{L}(\nu) = 1 - g(\nu)$ for $\nu > \bar{\nu}_a$. The cutoff $\underline{\nu}_a$ is the unique root in $(0, 1)$ of

$$\left(\frac{1}{a} - 1\right) g(\underline{\nu}_a)^2 - \left(\frac{1}{a} + y\right) g(\underline{\nu}_a) + x = 0, \tag{9}$$

$\bar{\nu}_a$ is the unique solution to $\mathcal{L}(\bar{\nu}_a) = 1 - g(\bar{\nu}_a)$, and $K_a$ is determined by $\mathcal{L}(\underline{\nu}_a) = g(\underline{\nu}_a)$.

(Note: The integral in (8), can be evaluated for all values of $a$ using the hypergeometric function. For particular values of $a$ simple evaluations can be obtained, for example when $a = y$ or when $a$ is an integer.)

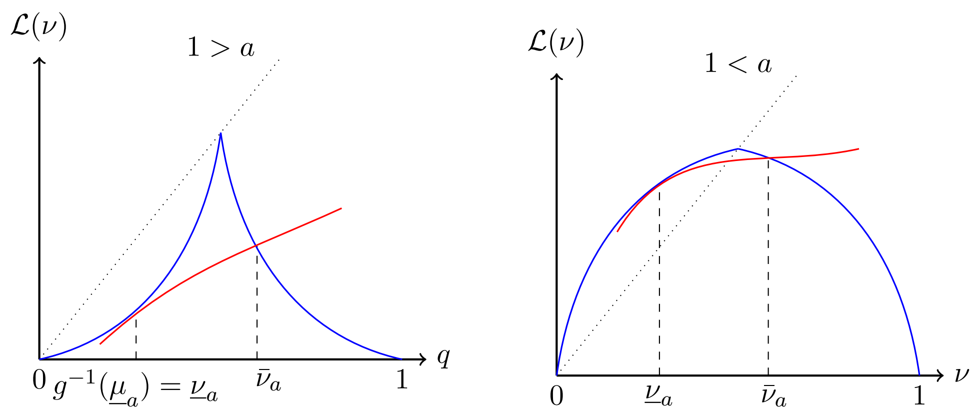


FIGURE 2. $\min\{g(\nu), 1 - g(\nu)\}$ (blue), the solution to the ODE (red), and the cutoffs

The cutoff $\underline{\mu}_a = g(\underline{\nu}_a)$, where sampling stops, can be evaluated for all values of $a$ using (9). For example, Bayesian updating, $a = 1$, generates the cutoff:

$$\underline{\mu}_1 = g(\underline{\nu}_1) = \frac{x}{1+y} = \frac{c}{\alpha}.$$

We can compare this Bayesian cutoff with the cutoff from general geometric updating. The cutoff belief, $g(\underline{\nu}_a)$, increases monotonically with $a$ and converges to zero as $a \to 0$.[24] As $a$ decreases, the geometric updating places increased weight on the new data that is sampled. From the agent's perspective the rate at which information arrives increases, although their costs of sampling are unchanged. Hence, they stop sampling only when their beliefs become more extreme. Conversely, agents who are less sensitive to new data stop sampling at less extreme beliefs.

[24] Implicit differentiation of (9) shows $dg(\underline{\nu}_a)/da^{-1} = -g(1-g)/(y + 2g + a^{-1}(1-2g)) < 0$.

Although, the agents with responsive updating stop sampling at extreme beliefs, they can achieve these extreme beliefs by collecting the same size sample as a Bayesian—it is not necessary for them to sample more. Nevertheless, agents who have responsive beliefs do actually sample more data than a Bayesian. To justify this, it is sufficient show that $\underline{\nu}_a$, the shadow cutoff, increases with $a$.[25] Thus a Bayesian observer of the agent's sampling behavior would have more extreme terminal beliefs as $a$ decreases.

### 5.2. *Sequential Sampling with Non-Divisible Updating*

Now we consider the same sequential sampling model, but with non-divisible updating. First, it is shown that if the agent updates continuously then the non-divisible updating is equivalent to divisible updating. Then it is shown that continuous updating may not be optimal, because the agent may learn more quickly if they commit to ignoring their signals for a predetermined time period. We end by discussing the optimal sampling when the agent can choose how frequently to update. We show that optimal sampling of a sophisticated agent generates some novel behavior.

The model of non-divisible updating in this section is a linear weighting of the prior and posterior similar to Epstein, Noor, and Sandroni (2010). It is not the same, because the weighting is not applied when a non-zero signal arrives.[26] To be precise, suppose the agent samples for $\Delta > 0$ periods and only then updates her beliefs. If her initial belief was $\mu_0$, then her updated beliefs are defined to be

$$\mu_\Delta = \begin{cases} 1 & \text{if } s_\tau = 1 \text{ for some } \tau \leq \Delta; \\ (1-\lambda)\mu_0 + \lambda \frac{\mu_0 e^{-\alpha\Delta}}{\mu_0 e^{-\alpha\Delta} + (1-\mu_0)e^{-\beta\Delta}} & \text{if } s_\tau = 0 \text{ for all } \tau \leq t; \\ 0 & \text{if } s_\tau = -1 \text{ for some, } \tau \leq \Delta. \end{cases} \tag{10}$$

The agent facing the sequential sampling problem with this non-divisible updating, potentially, has two choices to make at the belief $\mu_0$: First, as before, whether to cease sampling or not. Second, how large a sample should she collect before deciding whether to update—how large should $\Delta$ be?

We begin by considering an agent who finds it impossible to control $\Delta$ and, as time is continuous, she uses (10) to *continuously* update her beliefs.[27] In this case her sampling behavior is identical to that of an agent with divisible updating. In fact, her optimal policy

[25] Substitute $g(\nu) = \frac{\nu^{1/a}}{\nu^{1/a}+(1-\nu)^{1/a}}$ into (9), change variables to $\underline{r}^{1/a} := (\frac{\underline{\nu}}{1-\underline{\nu}})^{1/a}$, and then differentiate.

[26] We choose the updating (10), as it simplifies the solution to the sequential sampling problem. If beliefs jumped to $(1-\lambda)p_0$ not to zero, for example, it could be the case that the agent would choose to continue sampling after such a jump.

[27] An alternative explanation of this behavior is that the agent is unsophisticated and does not realize that her choice of $\Delta$ affects the updating.

and payoff is as described in Proposition 2 with $\lambda = 1/a$. Suppose that $\Delta = dt$ in (10), then the continuously updated belief process satisfies

$$\mu + d\mu = (1-\lambda)\mu + \lambda \frac{\mu(1-\alpha\, dt)}{\mu(1-\alpha\, dt) + (1-\mu)(1-\beta\, dt)}.$$

Thus $d\mu = \lambda(\beta - \alpha)\mu(1-\mu)\, dt$, or

$$\mu_t = \frac{\mu_0 e^{-\lambda\alpha t}}{\mu_0 e^{-\lambda\alpha t} + (1-\mu_0)e^{-\lambda\beta t}}.$$

This is identical to (6), when $f^{-1}(\mu) \equiv g(\mu) = \mu^\lambda/(\mu^\lambda + (1-\mu)^\lambda)$. Hence, the optimal sampling behavior when these non-divisible beliefs are continuously updated is identical to that described in Proposition 2 when $1/a = \lambda$: the agent samples for all $g^{-1}(\mu) \in (\underline{\nu}_{1/\lambda}, \bar{\nu}_{1/\lambda})$ as defined in (9).[28]

What if the agent can choose, $\Delta$, the period that she commits to ignore her signals without updating? Increasing $\Delta$ and committing to delay the updating can benefit her, because her beliefs change quicker. Lemma 3 compares $\mu_{2\Delta}$ (beliefs updated twice (11)), with $\hat{\mu}_{2\Delta}$ (the beliefs from one aggregate update (12)). The results, displayed in the figure below, show that there are regions where beliefs fall more quickly after only one aggregate update. But, there is no simple comparison between the beliefs from these two updating procedures. One might conjecture that when $\lambda < 1$ beliefs decline more rapidly if there are two updates rather than an aggregate update. This is not the case.

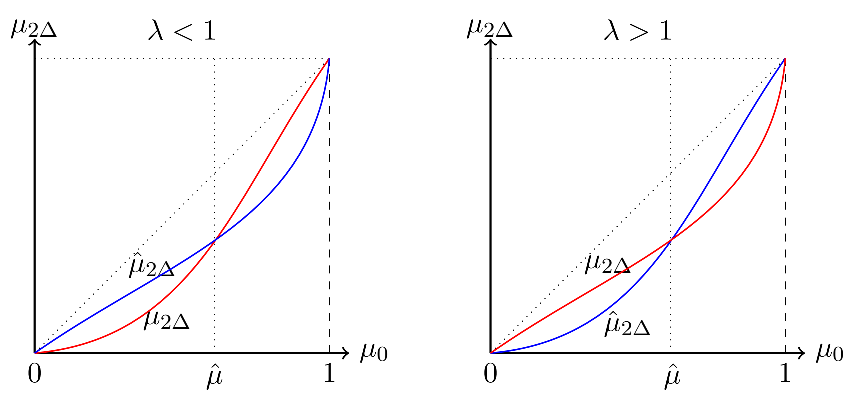


FIGURE 3. Speed of Learning: one update, $\hat{\mu}_{2\Delta}$ (blue), versus two $\mu_{2\Delta}$ (red).

[28]The equivalence here between divisible updating and continuously applied non-divisible updating in not unexpected. Once an updating procedure has been determined for the finest partition of time with this sort of signal structure. A divisible updating rule can be created from the continuously updated belief process by *defining* the updating over larger intervals to be the aggregation of the continuous updating. Here an agent who updates continuously (but non-divisibly) can be treated, *as if*, they were updating divisibly.

Before stating the Lemma we define $\hat{\mu}_{2\Delta}$ and $\mu_{2\Delta}$. First, the beliefs with two updates and only the zero signal

$$\mu_{2\Delta} := (1-\lambda)\mu_\Delta + \lambda\frac{\mu_\Delta}{\mu_\Delta + (1-\mu_\Delta)\zeta}, \qquad \mu_\Delta := (1-\lambda)\mu_0 + \lambda\frac{\mu_0}{\mu_0 + (1-\mu_0)\zeta}; \tag{11}$$

where $\zeta := e^{(\alpha-\beta)\Delta}$. Then, the beliefs with one aggregate update:

$$\hat{\mu}_{2\Delta} := (1-\lambda)\mu_0 + \lambda\frac{\mu_0}{\mu_0 + (1-\mu_0)\zeta^2}. \tag{12}$$

**Lemma 3.** *For each $\Delta$, $\lambda$ there exists a unique $\hat{\mu} \in (0,1)$ such that:*
(i) *If $\lambda < 1$, then: $\mu_{2\Delta} > \hat{\mu}_{2\Delta}$ for $\mu_0 < \hat{\mu}$; and $\mu_{2\Delta} < \hat{\mu}_{2\Delta}$ for $\mu_0 > \hat{\mu}$.*
(ii) *If $\lambda > 1$, then: $\mu_{2\Delta} < \hat{\mu}_{2\Delta}$ for $\mu_0 < \hat{\mu}$; and $\mu_{2\Delta} > \hat{\mu}_{2\Delta}$ for $\mu_0 > \hat{\mu}$.*
(iii) *$\hat{\mu} \to \frac{1}{2}$ as $\Delta \to 0$.*

Figure 3 suggests the optimal value of $\Delta$ will vary considerably. For example, when $\lambda > 1$ and $\mu_0 > \hat{\mu}$ the agent's learning will be fastest if she continuously updates. But as her beliefs fall, there may come a point where she prefers to stop continuous updating. At that point she prefers to commit to update at discrete times (some details on this are given in Lemma 4). Thus her beliefs will decline smoothly for a while and then proceed downwards in one or more jumps. When $\lambda < 1$ the reverse is true. The agent will first prefer to update the beliefs discretely and then switch to continuous belief revision for low values of $\mu$.

In the Lemma 4 we describe part of the optimal sampling policy when $\lambda > 1$. It shows that when $\mu_0$ is low, the optimal policy for the agent is to commit to not updating her beliefs for a predetermined period. This period will be so large that her eventual updated beliefs lie strictly inside the stopping region. The optimal policy is not a simple cutoff rule. First, it is shown that at the beliefs below the cutoff $\mu_0 \leq \underline{\mu}$ (defined by (9) with $a^{-1} = \lambda$) no update of any size is optimal: once $\mu \leq \underline{\mu}$ her updating ceases. However, there is an interval $\mu_0 \in (\underline{\mu}, \mu^\dagger)$ where the optimal policy is to commit to sample for a strictly positive time period before updating. The sampling period, $\Delta^*(\mu_0)$, is so large that should no (non-zero) signal arrive, her updated beliefs jump from $\mu_0 > \underline{\mu}$ to strictly below $\underline{\mu}$. This is the final inequality of the Lemma.

**Lemma 4.** *If: the updating is given by (10), $\lambda > 3(1+y)$, $\alpha > \beta > 2c$, and $\underline{\mu}$ is the unique solution to*

$$(\lambda - 1)\mu^2 - (y+\lambda)\mu + x = 0 \tag{13}$$

*in $(0, \frac{1}{2})$. If $\mu_0 \leq \underline{\mu}$ then no sampling is optimal. If $\mu_0 \in (\underline{\mu}, \mu^\dagger)$, for some $\mu^\dagger > \underline{\mu}$, then it is optimal to commit to sampling for a strictly positive time interval, $\Delta^*(\mu_0) > 0$, satisfying:*

$$(1-\lambda)\mu_0 + \lambda\frac{\mu_0 e^{-\alpha\Delta^*(\mu_0)}}{\mu_0 e^{-\alpha\Delta^*(\mu_0)} + (1-\mu_0)e^{-\beta\Delta^*(\mu_0)}} < \underline{\mu}.$$

# 6. Properties of Quasi-Bayesian Updating

In this section we will give conditions for the updating of Proposition 1 to be consistent and describe conditions on the bijection $F$ that generates desirable features.

## 6.1. The Consistency of Quasi-Bayesian Updating

In a stationary environment there is a property that might be desirable in a model of updating: the agent knows that their updated beliefs converge to the truth if they see enough data. This property is termed *consistency* when it holds in the Bayesian case, see Diaconis and Freedman (1986) for example. The result below shows that this property is also satisfied by quasi-Bayesian learning provided: the updater's model includes the true data generating process, the updating is continuous, and one weak further property. The consistency of this divisible updating contrasts with other examples of non-Bayesian updating in the literature that do not satisfy consistency.[29]

It is necessary to have a model where the agent repeatedly and independently samples from the same experiment $\mathcal{E}_n$ for a fixed state. Let an experiment $\mathcal{E}_n = (p^\theta)_{\theta\in\Theta}$ and a state $\tilde{\theta} \in \Theta$ be given. We use this to define three different stochastic processes. First, define the stochastic process $\{s^t\}_{t=0}^{\infty} \in S_n^{\infty}$ to be the independent and identically distributed signals that are sampled from the distribution $p^{\tilde{\theta}}$. Let $\mathbb{P}^{\tilde{\theta}}$ to denote the probability measure on $S_n^{\infty}$ induced by this process. Second, define $\{\mu^t\}_{t=0}^{\infty} \in \Delta(\Theta)^{\infty}$ to be the beliefs of an agent who observes the sequence of signals $\{s^t\}_{t=0}^{\infty} \in S^{\infty}$ and who updates according to Proposition 1. These are defined recursively, so that $\mu^{t+1}$ is the updated value of $\mu^t$ when the signal $s^t$ is observed, that is,

$$\mu^0 \in \Delta^o(\Theta), \qquad \mu^{t+1} := u\left(\mu^t, p_{s^t}\right), \qquad t = 0, 1, \dots \ . \tag{14}$$

Finally we define the stochastic process followed by the shadow beliefs $\{F(\mu^t)\}_{t=0}^{\infty}$. By Proposition 1 this process is updated using Bayes rule, so

$$F(\mu^{t+1}) = \frac{F(\mu^t) \circ p_{s_t}}{F(\mu^t)^T p_{s_t}}, \qquad t = 0, 1, \dots \ .$$

When the shadow prior gives positive probability to the true state ($F(\mu^0) \in \Delta^o(\Theta)$) and the signals can identify the state ($p^{\tilde{\theta}} \neq p^\theta$ for all $\theta \neq \tilde{\theta}$), the usual proof of the consistency of Bayesian updating (see for example DeGroot (1970)) applies to the shadow belief process. An immediate application of this tells us that $F(\mu^t) \to e_{\tilde{\theta}}$, $\mathbb{P}^{\tilde{\theta}}$ almost surely.[30]

[29] See Rabin and Schrag (1999) and Epstein, Noor, and Sandroni (2010) for examples of inconsistent updating. In Lehrer and Teper (2015) a notion of consistency is used as an axiom to characterize Bayesian updating.

[30] Here $e_{\tilde{\theta}}$ is the basis vector with unity in the $\tilde{\theta}$ entry and zeros elsewhere.

In order to show that quasi-Bayesian belief updating is consistent, therefore, we need to show that beliefs converge to the truth when the shadow beliefs converge to the truth. That is, we need to show that $F^{-1}(F(\mu^t)) \to e_{\tilde{\theta}}$ when $F(\mu^t) \to e_{\tilde{\theta}}$. This is clearly going to require continuity of $F^{-1}$ at the boundary. If $F^{-1}$ were not continuous there is no reason a limiting property continues to hold when $F^{-1}$ is applied to a convergent sequence. To ensure this we require $F$ to be continuous on its entire domain.[31] A second property of $F$ will also be necessary. The example of Bad Bayesian updating, of Section 4.3, is continuous. But in this case, $F^{-1}(e_{\tilde{\theta}}) \neq e_{\tilde{\theta}}$, so learning converges to a belief in the wrong state. To ensure quasi-Bayesian updating converges to the belief in the true state, we require that revealing signals are treated appropriately by the updating. To this end we will require that an agent who saw a signal that occurs with positive probability in state $\theta$ and with zero probability in all other states $\theta' \neq \theta$ correctly deduces the state:

$$u(\mu, e_\theta) = e_\theta, \qquad \forall\, \theta \in \Theta, \mu \in \Delta^o(\Theta). \qquad \text{(Respects Certainty)} \tag{15}$$

We will say that updating which satisfies (15) respects certainty. Equipped with these additional restrictions quasi-Bayesian updating will be consistent. This discussion is summarized in the result below.

**Result 1.** *Assume the updating $\mathcal{U}$ satisfies the Axioms 1–4, respects certainty and is continuous on $\Delta(\Theta) \times \Delta(S)^{|\Theta|}$. If $p^{\tilde{\theta}} \neq p^\theta$ for all $\theta \neq \tilde{\theta}$, then $\mu^t \to e_{\tilde{\theta}}$, $\mathbb{P}^{\tilde{\theta}}$ almost surely.*

### *6.2. Learning and Quasi-Bayesian Updating*

In this section we are interested in whether the expected value of the updated belief in $\theta$ (the true state) is greater than the original belief $\mu_\theta$. That is, we study whether updated beliefs on average move towards or away from the true value of $\theta$. We define two properties:

$$\mu_\theta \leq E^\theta(u_\theta(\mu, p_s)), \qquad \text{positive learning;}$$
$$\mu_\theta \geq E^\theta(u_\theta(\mu, p_s)), \qquad \text{negative learning.}$$

Positive learning says that an agent, with beliefs $\mu^\theta$ expects to have an increased belief in $\theta$ when they observe the outcome of an experiment generated by $\theta$. This holds globally for a Bayesian updater; it is the conditional submartingale property for Bayesian posteriors. Negative learning says that the belief in $\theta$ is expected to decrease when $\theta$ is true. This may happen as the agent is slow to move their belief in $\theta$ upwards in response to positive evidence, but quick to move beliefs down when evidence in favor of an alternative $\theta'$ is observed. It could be interpreted as a reluctance to move to extreme beliefs or a skeptical attitude to evidence. We will show that if the divisible updating satisfies the conditions of Proposition 1, then this negative learning cannot hold at all priors: it must be a local not a global property.

[31] Consider the example of discontinuous updating in Section 4.3 as $\mu \to 0$ so $F^{-1}(\mu) \to 0.5$. Thus $F^{-1}$ is discontinuous on the boundary because $F$ has an interior discontinuity.

Here we provide sufficient conditions for positive or negative learning in dichotomies, that is, experiments with two possible states $\Theta = \{\theta, \theta'\}$.[32] In dichotomies, such as the learning model of Section 5, the bijection $F : (\mu_\theta, \mu_{\theta'}) \mapsto (\mu'_\theta, \mu'_{\theta'})$ can be described by its effect on its first element: $F(\mu) \equiv (f(\mu_\theta), 1 - f(\mu_\theta))$, where $f : [0,1] \to [0,1]$ is a bijection. The quasi-Bayesian updating conditional on the signal $s$ can then be written explicitly as

$$u_\theta(\mu, p_s) = f^{-1}\left(\frac{f(\mu_\theta)p_s^\theta}{f(\mu_\theta)p_s^\theta + (1 - f(\mu_\theta))p_s^{\theta'}}\right). \tag{16}$$

Proposition 3 gives local conditions for positive or negative learning, so we define an appropriate neighborhood of the original belief $\mu_\theta$. Define the interval $R_f(\mu_\theta) \subset (0,1)$ so that it includes all possible realizations of the updated beliefs in $\theta$ when the prior is $\mu_\theta$.[33] The first of these conditions is that $f$ is increasing on $R_f(\mu_\theta)$. This ensures that evidence in favor of the state $\theta$ is interpreted as such and excludes the Bad Bayesian updating of Section 4.3. The second condition is convexity or concavity—this allows the mean preserving spread of the Bayesian updating of shadow beliefs to be translated by $F$ a more or less dispersed distribution of updated beliefs.

**Proposition 3.** Suppose the quasi-Bayesian updating in a dichotomy is described by (16). Then,

(i) If $f(.)$ is increasing and $\frac{1}{f(.)}$ convex on $R_f(\mu_\theta)$, then $\mu_\theta \le E^\theta(u_\theta(\mu, p_s))$.

(ii) If $f(.)$ is increasing and $\frac{1}{f(.)}$ concave on $R_f(\mu_\theta)$, then $\mu_\theta \ge E^\theta(u_\theta(\mu, p_s))$.

If $f(0) = 0$ then $\frac{1}{f(.)}$ is not concave on any interval of the form $(0, x)$.

For Bayesian updating $f(\mu) = \mu$ and $\frac{1}{\mu}$ is convex, so Proposition 3 confirms that Bayesian updating is a conditional submartingale. As another example, consider the geometric weighting of Section 4.1. For dichotomies this is described by the bijection $f(\mu) = \frac{\mu^a}{\mu^a + (1-\mu)^a}$. When $a \ge 1$, $\frac{1}{f(.)}$ is convex so the learning is a conditional submartingale like Bayesian updating. However, when $a < 1$ there is negative learning when $\mu > \frac{1}{2}(1 + a)$ as $\frac{1}{f(.)}$ is concave here. In this case, as $\mu$ approaches unity the agent's updated belief is expected to decrease. Nevertheless, the Result 1 shows that $\mu \to 1$ almost surely.

### *6.3. Sufficient Conditions for Under and Overreaction to Information*

In this section we give sufficient conditions for a quasi-Bayesian updating rule to overreact or under-react to new information. Overreaction has many meanings in the literature on updating. Here it is defined to hold if the log likelihood of the updated beliefs has a variance

[32] It is far from clear whether any result of this kind can be established when there are many states.

[33] Convex Hull$\{\mu_\theta, (u_\theta(\mu, p_s))_{s \in S}\} \subset R_f(\mu_\theta)$ .

that is greater than the variance of the log likelihood of a Bayesian's updated beliefs. The log-likelihood ratio of a Bayesian updater follows a homogeneous random walk. Thus its variance is a prior-independent measure of the variability of Bayesian updating. This can be seen from the simple calculation

$$\mathrm{Var}\left[\ln \frac{\hat{\mu}_\theta}{1-\hat{\mu}_\theta}\right] = \mathrm{Var}\left[\ln \frac{\mu_\theta}{1-\mu_\theta} + \ln \frac{p_s^\theta}{p_s^{\theta'}}\right] = \mathrm{Var}\left[\ln \frac{p^\theta}{p^{\theta'}}\right].$$

(Where $\hat{\mu}_\theta := \mu_\theta p_s^\theta / \mu^T p_s$ is the Bayesian update and the variance is taken unconditionally of the state value.) Thus $\mathrm{Var}\left[\ln \frac{p^\theta}{p^{\theta'}}\right]$ will be our benchmark and we will define under-reaction and overreaction relative to this as follows:

$$\mathrm{Var}\left[\ln \frac{\mathrm{u}_\theta(\mu, \mathrm{p_s})}{1-\mathrm{u}_\theta(\mu, \mathrm{p_s})}\right] > \mathrm{Var}\left[\ln \frac{\mathrm{p}^\theta}{\mathrm{p}^{\theta'}}\right], \qquad \forall \mu; \qquad \text{(overreaction)}$$

$$\mathrm{Var}\left[\ln \frac{\mathrm{u}_\theta(\mu, \mathrm{p_s})}{1-\mathrm{u}_\theta(\mu, \mathrm{p_s})}\right] < \mathrm{Var}\left[\ln \frac{\mathrm{p}^\theta}{\mathrm{p}^{\theta'}}\right], \qquad \forall \mu. \qquad \text{(under-reaction)}$$

There is a simple intuition for the sufficient conditions for overreaction given in Proposition 4. If the inverse homeomorphism $F^{-1}$ moves points apart, then when the shadow posteriors and prior are mapped back to the belief space they are even further apart. The response to the signals has become more exaggerated and overreaction is present. Similarly, if the function $F^{-1}$ is a contraction, then the learning that occurred in the shadow Bayesian world gets reduced when it is mapped back to the belief space by $F^{-1}$. As a result the Bayesian learning in the shadow space is understated and there is under-reaction to new information. In the one-dimensional case this means that the slope of the function $f(.)$ will play a role in characterizing under- or overreaction.

**Proposition 4.** Suppose that the divisible updating in a dichotomy, $u_\theta(\mu, p_s)$, is described by the function $f(.)$, as in (16), and that $f$ is continuously differentiable.
If $f'(\mu) > \frac{f(\mu)(1-f(\mu))}{\mu(1-\mu)}$ for all $\mu \in (0,1)$, then the updating exhibits under-reaction.
If $f'(\mu) < \frac{f(\mu)(1-f(\mu))}{\mu(1-\mu)}$ for all $\mu \in (0,1)$, then the updating exhibits overreaction.

As an example, we can apply this result to the geometric and exponential weighting above. Their associated bijections for dichotomies are: $f^a(\mu) := \frac{\mu^a}{\mu^a+(1-\mu)^a}$ and $f^b(\mu) := \frac{e^{-b/\mu}}{e^{-b/\mu}+e^{-b/(1-\mu)}}$.

$$(f^a)'(\mu) = a\frac{f^a(\mu)(1-f^a(\mu))}{\mu(1-\mu)}, \qquad (f^b)'(\mu) = f^b(\mu)(1-f^b(\mu))\left(\frac{b}{\mu^2}+\frac{b}{(1-\mu)^2}\right).$$

Thus geometric weighting exhibits under-reaction if $a > 1$ and overreaction if $a < 1$. Exponential weighting exhibits under-reaction if $b > 1/2$.

## 7. Characterization of Bayesian Updating

In this section we study updating that is unbiased (or a martingale, or satisfies Bayes plausibility). Two results are given: First we show that unbiased updating (with no further restrictions) is equivalent to Bayesian updating where the experiment is misspecified. Second, we characterize Bayesian updating when $|\Theta| > 2$. Updating is Bayesian if and only if it satisfies all of the conditions of Result 1 and one further condition that we call *"no learning without evidence"*.

We begin by motivating and defining the two main properties that are imposed on the updating in this section. Given an experiment $\mathcal{E}_n$, the updating $\mathcal{U}_n$ specifies a profile of $n$ updated beliefs $(\mathcal{U}_n^s)_{s\in S_n}$. We place restrictions on the relationship between these $n$ updates and the original belief $\mu$. The first restriction is that the initial belief, $\mu$, equals a particular weighted average of the $n$ updates. The weight on update $\mathcal{U}_n^s$ in this average is the ex-ante probability of the signal $s$. Thus the unconditional expectation of the agent's updated beliefs equals the original beliefs. This martingale property is satisfied by Bayesian updating and is also called Bayes plausibility.

**Axiom 5** (Unbiased)**.** For any $\mu \in \Delta^o(\Theta)$, $n > 1$, and $\mathcal{E}_n = (p^\theta)_{\theta\in\Theta} \in \mathfrak{E}$ the updating function $\mathcal{U}_n(\mu, \mathcal{E}_n) \equiv (\mathcal{U}_n^s)_{s\in S_n}$ satisfies $\mu \equiv \sum_{s\in S_n}(\mu^T p_s)\mathcal{U}_n^s$.

The second property we consider is much weaker. It considers the convex hull of the $n$ points $(\mathcal{U}_n^s)_{s\in S_n}$ and requires that this set contains the original belief $\mu$.

**Axiom 6** (No Learning without Evidence)**.** For any $\mu \in \Delta^o(\Theta)$, $n > 1$, and $\mathcal{E}_n = (p^\theta)_{\theta\in\Theta} \in \mathfrak{E}$ the updating function $\mathcal{U}_n(\mu, \mathcal{E}_n) \equiv (\mathcal{U}_n^s)_{s\in S_n}$ satisfies $\mu \in \text{Convex Hull}\left\{\mathcal{U}_n^1, \dots, \mathcal{U}_n^n\right\}$.

When this restriction is not satisfied the agent forming the updated beliefs knows for certain that, whatever signal/evidence they see, their beliefs, $\mu$, will move in a similar direction.[34] Hence, we say there is learning without evidence when Axiom 6 fails. An alternative interpretation of this failure is that there is some exogenous bias in the updating, so the beliefs move in a certain direction independently of the signal.

Now we show that any updating $\mathcal{U}$ satisfying Axiom 5 for the experiment $\mathcal{E}_n$ can be interpreted as application of Bayes rule to an alternative experiment $\tilde{\mathcal{E}}_n$. Unbiased updating can, therefore, be interpreted as Bayesian updating with a misspecified model. The experiment $\tilde{\mathcal{E}}_n$, in Proposition 5, has the same unconditional signal probabilities as $\mathcal{E}_n$. Thus the two experiments, $\mathcal{E}_n$ and $\tilde{\mathcal{E}}_n$, will agree on the empirical probabilities that signals occur. Furthermore, if $\mathcal{U}$ satisfies Axiom 5 it gives the same profile of updated beliefs as the Bayesian update for the experiment $\tilde{\mathcal{E}}_n$.

[34] The separating hyperplane theorem defines the direction that beliefs must move as they are updated from $\mu$ to the disjoint convex set of updates .

**Proposition 5.** The updating $\mathcal{U}$ satisfies Axiom 5, if and only if, for all $n > 1$, $\mu \in \Delta^o(\Theta)$ and $\mathcal{E}_n = (p_s)_{s\in S_n} \in \mathfrak{E}$ there exists $\tilde{\mathcal{E}}_n = (\tilde{p}_s)_{s\in S_n} \in \mathfrak{E}$ satisfying $\mu^T \tilde{p}_s = \mu^T p_s$ for all $s \in S_n$, such that $\mathcal{U}_n(\mu, \mathcal{E}_n) = \mathcal{U}^B_n(\mu, \tilde{\mathcal{E}}_n)$. Where

$$\mathcal{U}^B_n(\mu, \tilde{\mathcal{E}}_n) := \left(u^B(\mu, \tilde{p}_1), \ldots, u^B(\mu, \tilde{p}_n)\right), \qquad u^B(\mu, \tilde{p}_s) := \frac{\mu \circ \tilde{p}_s}{\mu^T \tilde{p}_s}.$$

We now consider what further properties need to be placed on quasi-Bayesian updating to ensure that it is Bayesian. One might view Proposition 1 as saying that divisibility updating is "almost" enough for Bayesian updating. As a consequence, there might be weak additional restrictions that, when added to Axioms 1–4, give Bayes updating. Axiom 6 is such restriction—it only requires that the initial beliefs are in the convex hull of the updates. Provided there are at least three different states, we show that this axiom has enough power to restrict divisible updating to be Bayesian. To be precise, when $|\Theta| \geq 3$ we show that updating satisfies Axioms 1–4 & 6, is continuous, and respects certainty, if and only if, the updating is Bayesian. The proof of Proposition 6 is built on considering experiments with only two signals. In this case Axiom 6 has greatest power, because it requires the initial belief to lie on a line that joins two updated beliefs (one for each signal). When the dimension of the belief space is large this imposes a significant restriction on the belief updates that can arise. Of course when the dimension of the belief space is also a line $|\Theta| = 2$ this Axiom has little effect.

**Proposition 6.** Suppose that $|\Theta| > 2$. $\mathcal{U}$ is continuous, satisfies (15), Axioms 1–4 and 6, if and only if,

$$\mathcal{U}_n(\mu, \mathcal{E}) = \left(u^B(\mu, p_1), \ldots, u^B(\mu, p_n)\right), \qquad u^B(\mu, p_s) := \frac{\mu \circ p_s}{\mu^T p_s}.$$

The proof of this result is long. The key step is to consider a binary experiment that reveals the state $\theta$ or not. By (15), the updated beliefs from such an experiment must include $e_\theta$ and one other point. The initial beliefs lie on the line joining these two points, by Axiom 6. We use this to argue that the homeomorphism $F$ maps line segments to line segments. And then that the ratios $F_{\theta'}(\mu)/F_{\theta''}(\mu)$ depend only on $\mu_{\theta'}$ and $\mu_{\theta''}$. Finally we solve a functional equation and show that $F$ is of the geometric weighting form with the weight $a = 1$.

## 8. Conclusion

We have introduced and characterized quasi-Bayesian updating a class of models of updating generated by the divisibility property. This updating satisfies desirable normative properties and generalizes Bayesian updating. This model of updating can be readily applied to dynamic models of costly information acquisition, such as sequential sampling. Using the characterization of quasi-Bayesian updating, we also provide an axiomatic foundation for Bayesian updating.

## References

Aczél, D., and M. Hosszú (1956): "On Transformations with Several Parameters and Operations in Multidimensional Spaces," *Acta Math. Acad. Sci. Hungar.*, 6, 327–338.

Angrisani, M., A. Guarino, P. Jehiel, and T. Kitagawa (2017): "Information Redundancy Neglect Versus Overconfidence: A Social Learning Experiment," Cemmap working paper, UCL.

Bardi, M., and I. Capuzzo-Dolcetta (2008): *Optimal control and viscosity solutions of Hamilton-Jacobi-Bellman equations*. Birkhäuser, Boston.

Bohren, A., and D. Hauser (2017): "Bounded Rationality and Learning, A Framework and a Robustness Result," *under review*, pp. 349–374.

Chauvin, K. P. (2019): "Euclidean Properties of Bayesian Updating," mimeo, Department of Economics, Harvard.

Che, Y.-K., and K. Mierendorff (2019): "Optimal dynamic allocation of attention," *American Economic Review*, 109(8), 2993–3029.

DeGroot, M. H. (1970): *Optimal Statistical Decisions*. McGraw-Hill, New York.

Diaconis, P., and D. Freedman (1986): "On the Consistency of Bayes Estimates," *The Annals of Statistics*, 14(1), 1–26.

Ely, J., A. Frankel, and E. Kamenica (2015): "Suspense and surprise," *Journal of Political Economy*, 123(1), 215–260.

Epstein, L. G., J. Noor, and A. Sandroni (2010): "Non-Bayesian Learning," *The B.E. Journal of Theoretical Economics*, 10(1).

Frankel, A., and E. Kamenica (2019): "Quantifying information and uncertainty," *American Economic Review*, 109(10), 3650–80.

Fudenberg, D., P. Strack, and T. Strzalecki (2018): "Speed, accuracy, and the optimal timing of choices," *American Economic Review*, 108(12), 3651–84.

Gilboa, I., and D. Schmeidler (1993): "Updating ambiguous beliefs," *Journal of economic theory*, 59(1), 33–49.

Grether, D. M. (1980): "Bayes rule as a descriptive model: The representativeness heuristic," *The Quarterly journal of economics*, 95(3), 537–557.

Hagmann, D., and G. Loewenstein (2017): "Persuasion with Motivated Beliefs," Carnegie Mellon University.

Hanany, E., and P. Klibanoff (2009): "Updating Ambiguity Averse Preferences," *The B.E. Journal of Theoretical Economics*, 9, 291–302.

Howson, C., and P. Urbach (2006): *Scientific reasoning: the Bayesian approach*. Open Court Publishing.

Kamenica, E., and M. Gentzkow (2011): "Bayesian persuasion," *American Economic Review*, 101(6), 2590–2615.

Legg, A. M., and K. Sweeny (2014): "Do you want the good news or the bad news first? The nature and consequences of news order preferences," *Personality and Social Psychology Bulletin*, 40(3), 279–288.

Lehrer, E., and R. Teper (2015): "Who is a Bayesian," Discussion paper, Working paper.

Levy, G., and R. Razin (2017): "Combining Forecasts: Why Decision Makers Neglect Correlation," Mimeo.

Lewis, D. (1976): "Probabilities of Conditionals and Conditional Probabilities," *The Philosophical Review*, 85(3), 297–315.

Majumdar, D. (2004): "An axiomatic characterization of Bayes' Rule," *Mathematical social sciences*, 47(3), 261–273.

Morris, S., and P. Strack (2017): "The Wald problem and the equivalence of sequential sampling and static information costs," *Unpublished manuscript, June*.

Nikandrova, A., and R. Pancs (2018): "Dynamic project selection," *Theoretical Economics*, 13(1), 115–143.

Ortoleva, P. (2012): "Modeling the change of paradigm: Non-Bayesian reactions to unexpected news," *American Economic Review*, 102(6), 2410–36.

Perea, A. (2009): "A model of minimal probabilistic belief revision," *Theory and Decision*, 67(2), 163–222.

Plott, C. R. (1973): "Path independence, rationality, and social choice," *Econometrica*, 41(6), 1075–1091.

Pomatto, L., P. Strack, and O. Tamuz (2018): "The cost of information," .

Rabin, M., and J. L. Schrag (1999): "First Impressions Matter: A Model of Confirmatory Bias," *The Quarterly Journal of Economics*, 114(1), 37–82.

Ramsey, F. P. (1926): "Truth and Probability," in *The Foundations of Mathematics and Other Logical Essays (1931)*, ed. by R. Braithwaite, chap. VII, pp. 156–198. London: Kegan, Paul,Trench, Trubner & Co.

Rubinstein, A., and L. Zhou (1999): "Choice problems with a 'reference' point," *Mathematical Social Sciences*, 37(3), 205–209.

Savage, L. J. (1972): *The Foundations of Statistics*. Dover, New York, NY, second revised edn.

Shannon, C. E., and W. Weaver (1949): *The Mathematical Theory of Communication*. University of Illinois Press, Urbana and Chicago.

Sokolov, M. V. (2011): "On scale-invariant parametric families of functions," *Journal of Mathematical Economics*, 47(4-5), 500–507.

Wald, A. (1945): "Sequential tests of statistical hypotheses," *The annals of mathematical statistics*, 16(2), 117–186.

Zhao, C. (2016): "Pseudo-Bayesian Updating," mimeo.

## Appendix

### *Proof of Lemma 1*

*Proof.* Axiom 3(a) says $\mathcal{U}_n^1(\mu, \mathcal{E}_n) \equiv \mathcal{U}_2^1(\mu, p_1, \mathbb{1} - p_1) := u(\mu, p_1)$. Symmetry, then, implies $\mathcal{U}_n^s(\mu, \mathcal{E}_n) \equiv u(\mu, p_s)$ for all $s$ and part (i) of the Lemma is established.

To establish part (iii) of this lemma rewrite the identity in Axiom 3(b) using (i). That is,

$$u(\mu, p_s) \equiv u\left(u\left(\mu, \mathbb{1} - p_1\right), p_s \circ (\mathbb{1} - p_1)^{-1}\right) \qquad \forall\ 0 < p_s \leq \mathbb{1} - p_1. \tag{17}$$

Defining $x = \mathbb{1} - p_1$ proves (iii).

To establish part (ii) take part(iii) and suppose that $x = \mathbb{1}\lambda^{-1}$ for $\lambda \in (1, \min_\theta(p_s^\theta)^{-1}]$. Axiom 1 applied to a binary experiment implies that the function $u(.)$ satisfies

$$u(\mu, p\mathbb{1}) = \mu, \qquad \forall \mu \in \Delta(\Theta), p \in (0,1). \tag{18}$$

This implies $u\left(\mu, x\right) \equiv \mu$. Writing part (iii) of the Lemma for this value of $x$ therefore gives

$$u(\mu, p_s) \equiv u\left(\mu, \lambda p_s\right), \qquad \forall \lambda \in [1, \min_\theta (p_s^\theta)^{-1}].$$

Hence the function $u(\mu, p)$ is homogeneous degree zero in $p$. □

### *Proof of Proposition 1*

*Proof.* By Lemma 1(i) we know that $\mathcal{U}_n(\mu, \mathcal{E}_n) \equiv (u(\mu, p_1), \ldots, u(\mu, p_n))$. We begin by transforming the variables in the function $u(\mu, p_s)$. Define $w : \Delta^o(\Theta) \to \mathbb{R}_{++}^{|\Theta|-1}$ as follows

$$w(\mu_1, \ldots, \mu_{|\Theta|}) := \left(\frac{\mu_1}{\mu_{|\Theta|}}, \ldots, \frac{\mu_{|\Theta|-1}}{\mu_{|\Theta|}}\right);$$

(where $\mathbb{R}_{++} := \{x \in \mathbb{R} : x > 0\}$). The function $w$ is a bijection, from $\Delta^o(\Theta)$ to $\mathbb{R}_{++}^{|\Theta|-1}$ and it has the inverse

$$w^{-1}(x_1, \ldots, x_{|\Theta|-1}) = \left(\frac{x_1}{1 + \sum_{i=1}^{|\Theta|-1} x_i}, \ldots, \frac{x_{|\Theta|-1}}{1 + \sum_{i=1}^{|\Theta|-1} x_i}, \frac{1}{1 + \sum_{i=1}^{|\Theta|-1} x_i}\right).$$

Redefine the variables of the function $u$ as: $\phi := w(\mu)$, $\pi := w(p_s)$ and $v(\phi, \pi) \equiv w(u(\mu, p_s))$. (The fact $u(\mu, p_s)$ that is homogeneous degree zero in $p_s \in (0,1)^{|\Theta|}$ (Lemma 1(ii)) implies there is no loss in this transformation.) The function $v : \mathbb{R}_{++}^{|\Theta|-1} \times \mathbb{R}_{++}^{|\Theta|-1} \to \mathbb{R}_{++}^{|\Theta|-1}$ satisfies (17), that is

$$v(\phi, \pi) \equiv v(v(\phi, \rho), \pi \circ \rho^{-1}), \qquad v : \mathbb{R}_{++}^{|\Theta|-1} \times \mathbb{R}_{++}^{|\Theta|-1} \to \mathbb{R}_{++}^{|\Theta|-1};$$

where: $\rho := w(\mathbb{1} - p_1)$, $\pi := w(p_s)$.

Now we do another transformation of the functional equation by taking logarithms,[35] that is, define $\tilde{\phi} := \ln \phi$, $\tilde{\pi} := \ln \pi$, $\tilde{\rho} = \ln \rho$, and $\tilde{v}(\tilde{\phi}, \tilde{\pi}) \equiv \ln v(\phi, \pi)$. A rewriting of (17) with this

[35] $\ln x$ where $x = (x_1, \ldots, x_n)$ denotes $(\ln x_1, \ldots, \ln x_n)$.

notation gives

$$\tilde{v}(\tilde{\phi},\tilde{\pi}) \equiv \tilde{v}(\tilde{v}(\tilde{\phi},\tilde{\rho}),\tilde{\pi}-\tilde{\rho}), \qquad\qquad \tilde{v}:\mathbb{R}^{|\Theta|-1}\times\mathbb{R}^{|\Theta|-1}\to\mathbb{R}^{|\Theta|-1}.$$

If we define $\tilde{y}=\tilde{\pi}-\tilde{\rho}$ and $\tilde{z}=\tilde{\rho}$ this then becomes the functional equation

$$\tilde{v}(\tilde{\phi},\tilde{y}+\tilde{z}) \equiv \tilde{v}\left(\tilde{v}(\tilde{\phi},\tilde{z}),\tilde{y}\right), \qquad \forall\tilde{\phi},\tilde{y},\tilde{z}\in\mathbb{R}^{|\Theta|-1}. \tag{19}$$

The functional equation (19) is called the translation equation. It was originally solved in its multivariate form by Aczél and Hosszú (1956). Given Axioms 4 and 2, their Theorem 2b (p. 331) applies here. Under the conditions assumed in the Proposition, $\tilde{v}$ solves (19) if and only if there exists a bijection $g:\mathbb{R}^{|\Theta|-1}\to\mathbb{R}^{|\Theta|-1}$ such that

$$\tilde{v}(\tilde{\phi},\tilde{\pi}) = g^{-1}[g(\tilde{\phi})+\tilde{\pi}]. \tag{20}$$

Now we will reverse the transformations of the variables used to derive (20). Substituting for the definitions of $\tilde{v}$, $\tilde{\phi}$, and $\tilde{\pi}$ gives

$$\ln v(\phi,\pi) = g^{-1}[g(\ln\phi)+\ln\pi].$$

Hence we have $g(\ln v(\phi,\pi)) = g(\ln\phi)+\ln\pi$. Now introduce the function $h(x) := g(\ln x) = (h_1(x),\ldots,h_{|\Theta|-1}(x))$ to simplify this expression.

$$\begin{aligned} h(\ v(\phi,\pi)\ ) &= h(\phi)+\ln\pi \\ e^{h(v(\phi,\pi))} &= \left(e^{h_1(\phi)}\pi_1,\ldots,e^{h_{|\Theta|-1}(\phi)}\pi_{|\Theta|-1}\right) \end{aligned}$$

We define $J(x)\equiv e^{h(x)}$, which allows the expression above to be rewritten as $J(v(\phi,\pi)) = \left(J_1(\phi)\pi_1,\ldots,J_{|\Theta|-1}(\phi)\pi_{|\Theta|-1}\right)$. Now substitute $v(.)=w(u(.))$ and $\pi=w(p_s)$ to get

$$J\left(w\left(u(\mu,p_s)\right)\right) = \left(J_1(w(\mu))\frac{p_s^1}{p_s^{|\Theta|}},\ldots,J_{|\Theta|-1}(w(\mu))\frac{p_s^{|\Theta|-1}}{p_s^{|\Theta|}}\right). \tag{21}$$

We will now define the function $F:\Delta^o(\Theta)\to\Delta^o(\Theta)$ so that the following diagram commutes, that is, $J(w(.))\equiv w(F(.))$. This is possible as $w$ is invertible and $J(.)=\exp(g(\ln(.)))$ is a bijection.

$$\begin{tikzcd} \Delta^o(\Theta) \arrow[r, "F"] \arrow[d, "w"] & \Delta^o(\Theta) \arrow[d, "w"] \\ \mathbb{R}^{|\Theta|-1}_{++} \arrow[r, "J"] & \mathbb{R}^{|\Theta|-1}_{++} \end{tikzcd}$$

$F$ is also bijection on $\Delta^o(\Theta)$. Using $J(w(.))=w(F(.))$ we can rewrite (21) as

$$\begin{aligned} w\left(F\left(u(\mu,p_s)\right)\right) &= \left(w_1(F(\mu))\frac{p_s^1}{p_s^{|\Theta|}},\ldots,w_{|\Theta|-1}(F(\mu))\frac{p_s^{|\Theta|-1}}{p_s^{|\Theta|}}\right), \\ &= \left(\frac{F_1(\mu)p_s^1}{F_{|\Theta|}(\mu)p_s^{|\Theta|}},\ldots,\frac{F_{|\Theta|-1}(\mu)p_s^{|\Theta|-1}}{F_{|\Theta|}(\mu)p_s^{|\Theta|}}\right). \end{aligned}$$

Now applying $w^{-1}$ to both sides gives

$$F\left(u\left(\mu,p_s\right)\right) \equiv \left(\frac{F_1(\mu)p_s^1}{F(\mu)^T p_s},\ldots,\frac{F_{|\Theta|}(\mu)p_s^{|\Theta|}}{F(\mu)^T p_s}\right). \tag{22}$$

Applying $F^{-1}$ to both sides of this gives (4). The other displayed equation in the proposition follows from a Lemma 1.

The above completes one direction of the proof. Now suppose that

$$\mathcal{U}_n(\mu,\mathcal{E}) \equiv (u(\mu,p_1),\dots,u(\mu,p_n))\,, \qquad \text{where} \qquad u(\mu,p_s) \equiv F^{-1}\!\left(\frac{F(\mu)\circ p_s}{F(\mu)^T p_s}\right);$$

for some bijection $F:\Delta^o(\Theta)\to\Delta^o(\Theta)$. We must show that this updating satisfies our axioms. If the experiment is uninformative then $p_s = k_s\mathbb{1}$ for some constant $k_s$, so

$$u(\mu,p_s) \equiv F^{-1}\!\left(\frac{k_s F(\mu)}{k_s}\right) = \mu.$$

Thus Axiom 1 is satisfied. Suppose the signal names are permuted then

$$\mathcal{U}_n(\mu,(\omega(p^\theta))_{\theta\in\Theta}) = \big(u(\mu,p_{\omega(1)}),\dots,u(\mu,p_{\omega(n)})\big) = \Big(\mathcal{U}_n^{\omega(1)}(\mu,\mathcal{E}_n),\dots,\mathcal{U}_n^{\omega(n)}(\mu,\mathcal{E}_n)\Big)$$

and Axiom 2 is satisfied. To verify that Axiom 3 is satisfied it is sufficient to verify that $u(\mu,p_s)\equiv u\left(u\left(\mu,x\right),p_s\circ x^{-1}\right)$. This follows from the following calculation

$$\begin{aligned}
u\left(u\left(\mu,x\right),p_s\circ x^{-1}\right) &= F^{-1}\!\left(\frac{F(u\left(\mu,x\right))\circ p_s\circ x^{-1}}{F(u\left(\mu,x\right))^T(p_s\circ x^{-1})}\right)\\
&= F^{-1}\!\left(\frac{\frac{F(\mu)\circ x}{F(\mu)^T x}\circ p_s\circ x^{-1}}{\left(\frac{F(\mu)\circ x}{F(\mu)^T x}\right)^T(p_s\circ x^{-1})}\right)\\
&= F^{-1}\!\left(\frac{F(\mu)\circ p_s}{\left(F(\mu)\right)^T p_s}\right)\equiv u(\mu,p_s)
\end{aligned}$$

We finally need to show that for some $\mu^o$ the equation $u(\mu^o,.)=\mu$ has a unique solution $p\in\Delta^o(\Theta)$. This is to verify that Axiom 4 is satisfied. As $F$ is a bijection we can choose $\mu^o$ so that $F(\mu^o)=|\Theta|^{-1}\mathbb{1}$, thus (as $\mathbb{1}^T p=1$)

$$u(\mu^o,p) = F^{-1}\!\left(\frac{F(\mu^o)\circ p}{F(\mu^o)^T p}\right) = F^{-1}(p).$$

As $F$ and $F^{-1}$ are both bijections, the equation $F^{-1}(.)=\mu$ has a unique solution for all $\mu$. □

*Proof of Lemma 2*

*Proof.* The “if’ part of the proof has already been verified. So, suppose that

$$u^j(\mu,p_s) = \frac{\mu\circ J(p_s)}{\mu^T J(p_s)} \equiv F^{-1}\left(\frac{F(\mu)\circ p_s}{F(\mu)^T p_s}\right)$$

for some homeomorphism $F$. Choose $p_s=\mathbb{1}x$ for some $x>0$. In this case, the above can be re-written as

$$\big(\mu_1 j_1(x),\dots,\mu_{|\Theta|}j_{|\Theta|}(x)\big) \equiv (\mu^T J(\mathbb{1}x))\mu.$$

This implies $j_\theta(x) = j_{\theta'}(x)$ for all $\theta, \theta' \in \Theta$ and all $x$. Thus the functions in the weighting are all identical and henceforth will be written as $j(.)$. As $F$ is a homeomorphism there exists $\hat{\mu}$ so that $F(\hat{\mu}) = \mathbb{1}(|\Theta|)^{-1}$. Substitute this value into the above identity

$$\frac{\hat{\mu} \circ J(p_s)}{\hat{\mu}^T J(p_s)} \equiv F^{-1}\left(\frac{p_s}{\mathbb{1}^T p_s}\right).$$

Now take the ratio of two entries in the vectors on each side of this expression

$$\frac{\hat{\mu}_\theta j(p_s^\theta)}{\hat{\mu}_{\theta'} j(p_s^{\theta'})} \equiv \frac{F_\theta^{-1}\left(\frac{p_s}{\mathbb{1}^T p_s}\right)}{F_{\theta'}^{-1}\left(\frac{p_s}{\mathbb{1}^T p_s}\right)}$$

If the vector $p_s$ is multiplied by the scalar $z$ RHS of the above is unaltered, thus the ratio on the left is homogeneous degree zero, so $\frac{\hat{\mu}_\theta j(p_s^\theta)}{\hat{\mu}_{\theta'} j(p_s^{\theta'})} \equiv \frac{\hat{\mu}_\theta j(zp_s^\theta)}{\hat{\mu}_{\theta'} j(zp_s^{\theta'})}$. Rewriting this gives

$$\frac{j(zp_s^{\theta'})}{j(p_s^{\theta'})} \equiv \frac{j(zp_s^\theta)}{j(p_s^\theta)}.$$

Letting $\phi := \ln j$ $x := \ln p_s^{\theta'}$ $x + y := \ln p_s^\theta$ and $w = \ln z$ this becomes

$$\phi(x+w) - \phi(x) = \phi(x+y+w) - \phi(x+y)$$

Or for a fixed value of $x$

$$\overbrace{\phi(w+x) - \phi(x)}^{k(w)} + \overbrace{\phi(y+x) - \phi(x)}^{k(y)} = \overbrace{\phi(y+w+x) - \phi(x)}^{k(w+y)}.$$

This is Cauchy's functional equation and has the solution $k(w) = Cw$ for some constant $C$. Thus $\phi(w) = Cw + D$ for constants $C$ and $D$, or $\ln j(z) = C \ln z + D$. This finally gives us $j_\theta(z) = j(z) = Kz^a$ where $a = C$ for all $\theta$. Substituting this into the expression for the updating gives

$$u^j(\mu, p_s) = \frac{\mu \circ J(p_s)}{\mu^T J(p_s)} = \frac{\mu \circ (p_s)^a}{\mu^T (p_s)^a},$$

which proves the claim. □

*Proof of Proposition 2*

*Proof.* We begin by describing a candidate solution to the HJB equation. Then we show that this is a viscosity solution to the HJB and appeal to Bardi and Capuzzo-Dolcetta (2008) Theorem III.4.11 to verify that this is the value function.

First, it is shown that it is optimal to sample when $\nu = \frac{1}{2}$ for all $a > 0$, hence $\underline{\nu}_a < \frac{1}{2} < \bar{\nu}_a$. At $\nu \leq \frac{1}{2}$ the loss from immediate stopping is $g(\nu)$. The loss from $dt$ periods more sampling at $\nu$ is on the right below. Thus, it is optimal to do at least $dt$ periods more sampling if

$$g(\nu) > c\,dt + g(\nu + d\nu)[1 - \alpha g(\nu)dt - \beta(1 - g(\nu))dt].$$

Dividing by $dt$ and recalling the above expression for $d\nu$ this inequality becomes

$$g'(\nu)\nu(1-\nu) > x - g(\nu)[y + g(\nu)], \qquad y = \frac{\beta}{\alpha - \beta}, x = \frac{c}{\alpha - \beta}.$$

As $g(\nu) = \nu^{\frac{1}{a}}/(\nu^{\frac{1}{a}} + (1-\nu)^{\frac{1}{a}})$ and $g'(\nu)\nu(1-\nu) = \frac{1}{a}g(\nu)(1-g(\nu))$, this is equivalent to

$$0 > \left(\frac{1}{a} - 1\right) g(\nu)^2 - \left(y + \frac{1}{a}\right) g(\nu) + x.$$

When $g(\frac{1}{2}) = \frac{1}{2}$ and $\alpha > \beta > 2c$, the inequality above holds for all $a > 0$. This shows it is optimal to sample when $\nu = \frac{1}{2}$, which implies $\underline{\nu}_a < 1/2$ and the loss is $g(\underline{\nu}_a)$ when sampling stops.

Rewriting (7) using the variables $r := \frac{\nu}{1-\nu}$ and $L(r) \equiv \mathcal{L}(\nu)$ gives:

$$rL'(r) + \left(g\left(\frac{r}{1+r}\right) + y\right) L(r) = x.$$

Substituting the specific form of $g$ gives the ODE

$$L'(r) + \left(\frac{r^{\frac{1}{a}-1}}{1 + r^{\frac{1}{a}}} + \frac{y}{r}\right) L(r) = \frac{x}{r}.$$

After multiplying through by the factor $r^y(1 + r^{\frac{1}{a}})^a$, both sides of this can be integrated in the usual way. A rearranging, then, gives

$$L(r) = \frac{x \int r^{y-1}(1 + r^{1/a})^a \, dr + K}{r^y(1 + r^{1/a})^a}. \tag{23}$$

We now find a solution to (23) that "smooth pastes" to $g(\nu)$ at $\underline{\nu}_a$: $L^*(\underline{\nu}_a) = g(\underline{\nu}_a)$ and $L^*(\underline{\nu}_a) = g'(\underline{\nu}_a)$ where $L^*(\nu) \equiv L(\frac{\nu}{1-\nu}) = L(r)$. This ensures the agent is indifferent between immediate stopping and $dt$ periods more sampling. If these two conditions are substituted into (7) evaluated at $\underline{\nu}_a$ and a further substitution of $g'(\underline{\nu}_a)\underline{\nu}_a(1-\underline{\nu}_a) = \frac{1}{a}g(\underline{\nu}_a)(1-g(\underline{\nu}_a))$ is performed we get the condition (9) that defines $\underline{\nu}_a$. ($\alpha > \beta > 2c$ ensures (9) has one solution in $(0, \frac{1}{2})$.) Thus we choose the solution to the ODE

$$L(r) = \frac{x \int_{\underline{r}_a}^{r} \rho^{y-1}(1 + \rho^{1/a})^a \, d\rho + K_a}{r^y(1 + r^{1/a})^a}. \tag{24}$$

$K_a$ is chosen so that $L^*(\underline{\nu}_a) = L(\frac{\underline{\nu}_a}{1-\underline{\nu}_a}) = g(\underline{\nu}_a)$, when $\underline{r}_a := \frac{\underline{\nu}_a}{1-\underline{\nu}_a}$.

We now show that for $\nu \in [\underline{\nu}_a, 1]$ the function $L(\frac{\nu}{1-\nu})$, in (24) has a unique intersection with $\min\{g(\nu), 1 - g(\nu)\}$ from below. We first show that $L(\frac{\nu}{1-\nu})$ does not intersect $g(\nu)$ at any $\nu > \underline{\nu}_a$. A sufficient condition for $g(\nu) \geq L(\frac{\nu}{1-\nu})$ is

$$r^{1/a} r^y (1 + r^{1/a})^{a-1} - x \int_{\underline{r}_a}^{r} \rho^{y-1}(1 + \rho^{1/a})^a \, d\rho - K_a \geq 0, \qquad \forall r \geq \underline{r}_a.$$

The derivative of this is

$$r^{y-1}(1 + r^{1/a})^{a-2} \left\{ r^{2/a}(1 + y - x) + r^{1/a}(y - 2x + a^{-1}) - x \right\}.$$

The term in braces is zero at $\underline{\nu}_a$ and increases thereafter. Hence the above difference is positive for all $\nu > \underline{\nu}_a$. Now we show that for $\nu \in [\underline{\nu}_a, 1]$ the function $L^*(\nu) = L(\frac{\nu}{1-\nu})$ is increasing and hence has a unique intersection with the decreasing function $1 - g(\nu)$. Rearranging (7) gives

$$(\alpha - \beta)^{-1} \nu(1-\nu) \mathcal{L}'(\nu) = \mathcal{L}(\nu)(y + g) - x.$$

Thus a sufficient condition for $L$ to be increasing is $L(\frac{\nu}{1-\nu})(y+g)-x>0$. As $L(\frac{\nu}{1-\nu}) \geq L(\frac{\underline{\nu}_a}{1-\underline{\nu}_a})$ and $g$ is increasing

$$L(\frac{\nu}{1-\nu})(y+g)-x \geq \mathcal{L}(\underline{\nu}_a)(y+g(\underline{\nu}_a))-x = g(\underline{\nu}_a)(y+g(\underline{\nu}_a))-x = \frac{g^2(\underline{\nu}_a)+g(\underline{\nu}_a)}{a} > 0.$$

(The first equality substitutes $\mathcal{L}(\underline{\nu}_a)=g$ and the second substitutes from (9) the condition defining $\underline{\nu}_a$.) Hence the sufficient condition holds, $L$ is increasing, and there is a unique $\bar{\nu}_a$ where

$$L\left(\frac{\bar{\nu}_a}{1-\bar{\nu}_a}\right) = 1-g\left(\frac{\bar{\nu}_a}{1-\bar{\nu}_a}\right).$$

We can now define our candidate solution to the HJB equation as $L^*(\nu) := L(\frac{\nu}{1-\nu})$ where

$$L(r) = \begin{cases} g\left(\frac{r}{1+r}\right) & \frac{r}{1+r} \leq \underline{\nu}_a; \\ \frac{x\int_{\underline{r}_a}^{r} \rho^{y-1}(1+\rho^{1/a})^a \, d\rho + K_a}{r^y(1+r^{1/a})^a} & \underline{\nu}_a < \frac{r}{1+r} < \bar{\nu}_a; \\ 1-g\left(\frac{r}{1+r}\right) & \frac{r}{1+r} \geq \bar{\nu}_a. \end{cases} \tag{25}$$

By Bardi and Capuzzo-Dolcetta (2008) Theorem III.4.11 (suitably adapted to allow state dependent discounting) the value function $\mathcal{L}(\nu)$ is the unique viscosity solution to the HJB. Thus we need to show that $L^*(\nu)$ is a viscosity solution to the HJB to verify that it is the value function. When it is differentiable $L^*(\nu)$ was constructed to be a solution to the HJB. It is not differentiable at one point, $\bar{\nu}_a$ where it has an upward kink. So, we must show that $L^*(\nu)$ is a viscosity subsolution to the HJB equation at $\bar{\nu}_a$. That is,

$$0 \geq w(\alpha-\beta)\bar{\nu}_a(1-\bar{\nu}_a) - c + L^*(\bar{\nu}_a)[\alpha g(\bar{\nu}_a) + \beta(1-g(\bar{\nu}_a))]$$

for all superdifferentials $w$ of $L^*$. This is an upwards intersection, so an upper bound on the superdifferential is given by the slope of the function below $\bar{\nu}_a$ that is

$$w \leq \lim_{\nu \to \bar{\nu}_a^-} (L^*)'(\bar{\nu}_a) = \frac{x}{\bar{\nu}_a(1-\bar{\nu}_a)} - \frac{g(\bar{\nu}_a)+y}{\bar{\nu}_a(1-\bar{\nu}_a)} L^*(\bar{\nu}_a)$$

a substitution of this upper bound verifies the subsolution condition and we can therefore deduce that $L^*(\nu)$ is the value function. This is exactly the function denoted $\mathcal{L}(\nu)$ in the Proposition. □

## *Proof of Lemma 3*

*Proof.* We will decompose the difference $\mu_{2\Delta} - \hat{\mu}_{2\Delta}$ into three terms $(\mu_{2\Delta}-\mu_\Delta)+(\mu_\Delta-\mu_0)-(\hat{\mu}_{2\Delta}-\mu_0)$. The last two terms in this decomposition equal

$$\begin{aligned} \frac{\mu_\Delta - \mu_0 - (\hat{\mu}_{2\Delta}-\mu_0)}{\lambda} &= \frac{\mu_0(1-\mu_0)(1-\zeta)}{\mu_0+(1-\mu_0)\zeta} - \frac{\mu_0(1-\mu_0)(1-\zeta^2)}{\mu_0+(1-\mu_0)\zeta^2} \\ &= \frac{-\mu_0(1-\mu_0)(1-\zeta)\zeta}{AC}. \end{aligned}$$

Where $A := \mu_0+(1-\mu_0)\zeta$ and $C = \mu_0+(1-\mu_0)\zeta^2$. Now adding the first term

$$\frac{\mu_{2\Delta}-\hat{\mu}_{2\Delta}}{\lambda(1-\zeta)} = \frac{\mu_\Delta(1-\mu_\Delta)}{B} - \frac{\mu_0(1-\mu_0)\zeta}{AC}$$

where $B := \mu_\Delta + (1-\mu_\Delta)\zeta$. This can be rewritten as

$$\frac{\mu_{2\Delta}-\hat{\mu}_{2\Delta}}{\lambda(1-\zeta)} = \frac{\mu_0^2}{ABC}\left((1-\mu_\Delta)\mu_\Delta\frac{AC}{\mu_0^2} - \frac{1-\mu_0}{\mu_0}\zeta B\right).$$

Writing the expression in brackets in terms of odds ratios permits a further factorization.

$$\frac{\mu_{2\Delta}-\hat{\mu}_{2\Delta}}{\lambda(1-\zeta)} = \frac{\mu_0^2\mu_\Delta^2}{ABC}\left(\frac{1-\mu_\Delta}{\mu_\Delta} - \zeta\frac{1-\mu_0}{\mu_0}\right)\left(1 - \frac{1-\mu_0}{\mu_0}\frac{1-\mu_\Delta}{\mu_\Delta}\zeta^2\right)$$

Now a substitution for $\mu_\Delta$ in the parentheses gives

$$\frac{\mu_{2\Delta}-\hat{\mu}_{2\Delta}}{\lambda(1-\zeta)^2} = \frac{(1-\lambda)(1-\mu_0)\mu_0\mu_\Delta^2}{(\lambda+(1-\lambda)A)BC}\left(1-\zeta^2\left(\frac{1-\mu_0}{\mu_0}\right)^2\left(1+\frac{\lambda(\zeta-1)}{\lambda+(1-\lambda)A}\right)\right).$$

Define $\hat{\mu}$ as the solution to the equation

$$1 = \zeta^2\left(\frac{1-\hat{\mu}}{\hat{\mu}}\right)^2\left(1+\frac{\lambda(\zeta-1)}{\lambda+(1-\lambda)(\hat{\mu}+(1-\hat{\mu})\zeta)}\right).$$

(The right is decreasing in $\hat{\mu}$ for all $\zeta > 0$, so $\hat{\mu}$ is uniquely defined and $\hat{\mu} = \frac{1}{2}$ when $\zeta = 1$.) It is clear from the final equality that the sign of $\mu_{2\Delta} - \hat{\mu}_{2\Delta}$ is determined by whether $\lambda > 1$ and whether $\mu > \hat{\mu}$ as is claimed in the Lemma. □

## *Proof of Lemma 4*

*Proof.* Consider an agent committing to $\Delta$ periods of sampling before updating their beliefs. We will suppose that they irrevocably cease experimentation should no revealing signal appear and consequently experience the loss $\mu_\Delta$. This is the minimization problem

$$\min_{\Delta\geq 0} c\Delta + (\mu e^{-\alpha\Delta} + (1-\mu)e^{-\beta\Delta})\mu_\Delta,$$

where $\mu_\Delta$ is given by (10). Making this substitution

$$\min_{\Delta\geq 0} c\Delta + \mu e^{-\alpha\Delta} + (1-\lambda)\mu(1-\mu)(e^{-\beta\Delta} - e^{-\alpha\Delta}). \tag{26}$$

The slope of this function divided by $(\alpha-\beta)e^{-\alpha\Delta} > 0$ is

$$\underbrace{\{(\lambda-1)\mu^2 - (\lambda+y)\mu + x\}}_{(13)} + \underbrace{\left[x(e^{\alpha\Delta}-1) + (\lambda-1)\mu(1-\mu)y(e^{(\alpha-\beta)\Delta}-1)\right]}_{\geq 0}. \tag{27}$$

Consider the two parts of (27). The first part is independent of $\Delta$. The final part is positive and increases in $\Delta$. So as $\Delta$ increases, the slope is either always positive or first negative and then positive. The first part of (27) is (13), so if $\mu = \underline{\mu}$ it is zero. For $\mu < \underline{\mu}$, the first term is strictly positive. Hence, (27) is strictly positive for all $\Delta$ and $\Delta = 0$ solves the minimization. For $\mu = \underline{\mu}$, the first term is zero. In this case $\Delta = 0$ is a stationary point and solves the minimization. In summary, for initial beliefs $\mu \leq \underline{\mu}$ no further experimentation is optimal.

If $\mu > \underline{\mu}$, the first term in (27) is negative and there is a unique strictly positive stationary point, $\Delta^*(\mu)$, that is a global minimum. We now show that $\Delta^*(\mu)$ strictly increases in $\mu$; thus higher initial beliefs give rise to a larger optimal update. For $\mu > \underline{\mu}$, $\Delta^*(\mu)$ makes (27) equal zero, so implicit differentiation then gives

$$\frac{d\Delta^*(\mu)}{d\mu} = \frac{(y+1)(1+(\lambda-1)(1-2\mu)) - (\lambda-1)(1-2\mu)ye^{(\alpha-\beta)\Delta}}{x\alpha e^{\alpha\Delta} + \beta(\lambda-1)\mu(1-\mu)e^{(\alpha-\beta)\Delta}}.$$

Substituting for $e^{(\alpha-\beta)\Delta}$ from (27) equated to zero, this can be rewritten as

$$\frac{d\Delta^*(\mu)}{d\mu} = \frac{(y+1)\frac{\mu}{1-\mu} + \frac{1-2\mu}{\mu(1-\mu)}xe^{\alpha\Delta}}{\beta x e^{\alpha\Delta} + \alpha\mu + \alpha(\lambda-1)\mu(1-\mu)} > 0. \tag{28}$$

When the agent has completed $\Delta^*(\mu)$ periods of sampling and no signal has arrived, she will update her beliefs to $\mu_{\Delta^*(\mu)}$, as defined by (10). We now show that for $\mu$ sufficiently close to $\underline{\mu}$, these updated beliefs are strictly below $\underline{\mu}$, that is, the agent chooses $\Delta^*(\mu)$ so their updated beliefs jump into the region $\mu < \underline{\mu}$. First, we calculate how the updated beliefs vary with $\mu$. Differentiating $\mu_{\Delta^*(\mu)}$ in (10) with respect to $\mu$ gives

$$\frac{d\mu_{\Delta^*(\mu)}}{d\mu} = 1 - \lambda + \lambda\frac{\mu^b(1-\mu^b)}{\mu(1-\mu)}\left(1 - (\alpha-\beta)\mu(1-\mu)\frac{d\Delta^*(\mu)}{d\mu}\right),$$

where $\mu^b = (1+\frac{1-\mu}{\mu}e^{(\alpha-\beta)\Delta})^{-1}$ is the Bayesian update. Substituting from (28) and evaluating this at $\underline{\mu}$ where $\Delta^*(\underline{\mu}) = 0$ we get

$$\left.\frac{d\mu_{\Delta^*(\mu)}}{d\mu}\right|_{\mu=\underline{\mu}} = 1 - \lambda\frac{\alpha\underline{\mu}^2 + (1-2\underline{\mu})c}{\beta x + \lambda\alpha\underline{\mu} - \alpha(\lambda-1)\underline{\mu}^2}.$$

As $\underline{\mu} = \mu_{\Delta^*(\underline{\mu})}$, if this derivative is negative then $\underline{\mu} > \mu_{\Delta^*(\mu)}$ for an interval of $\mu > \underline{\mu}$ and the updated beliefs fall below this threshold.

This derivative is negative if $\alpha(2\lambda-1)\underline{\mu}^2 - \lambda(\alpha+2c)\underline{\mu} + c(\lambda-y) > 0$. Re-arranging this inequality and writing the definition of $\underline{\mu}$ in a similar way gives:

$$A + B\underline{\mu} := \frac{c(\lambda-y) + \underline{\mu}(\alpha(\lambda-1) - 2c\lambda)}{\alpha(2\lambda-1)} > \underline{\mu}(1-\underline{\mu}),$$

$$C + D\underline{\mu} := \frac{x-(y+1)\underline{\mu}}{\lambda-1} = \underline{\mu}(1-\underline{\mu}).$$

We now provide sufficient conditions for $A > C$ and $B > D$ (when $\lambda > 1$), so the inequality holds for all positive $\underline{\mu}$. First, note that $A > C$ iff $0 < \lambda^2 - 3\lambda(1+y) + 1 + 2y$; a sufficient condition this is our assumption that $\lambda > 3(1+y)$. Second, note that $B > 0 > D$ if $\lambda > \frac{\alpha}{\alpha-2c}$, this is true if $\lambda > 3(1+y)$ and $\beta > 2c$. □

## *Proof of Proposition 3*

*Proof.* We begin by establishing part (i) of the result. Re-arranging (16) we get a relation for $u_\theta(\mu, p^s)$, the updated belief in the state $\theta$ conditional on the signal $s$,

$$f(u_\theta(\mu, p^s)) = \frac{f(\mu_\theta)p_s^\theta}{f(\mu_\theta)p_s^\theta + (1-f(\mu_\theta))p_s^{\theta'}}.$$

This is the shadow posterior on $\theta$. Now we calculate the shadow posterior odds ratio. Then take its expectation conditional on the state $\theta$, that is

$$E^\theta\left(\frac{1-f(u_\theta(\mu,p^s))}{f(u_\theta(\mu,p^s))}\right)=\sum_s p_s^\theta\frac{(1-f(\mu_\theta))p_s^{\theta'}}{f(\mu_\theta)p_s^\theta}=\frac{1-f(\mu_\theta)}{f(\mu_\theta)}.$$

Adding unity to the extremes of this equality gives

$$E^\theta\left(\frac{1}{f(u_\theta(\mu,p^s))}\right)=\frac{1}{f(\mu_\theta)}.$$

The function $\frac{1}{f(.)}$ is assumed to be convex on an interval of values containing the points $\{u_\theta(\mu,p^s): s\in S\}$. Therefore, by Jensen's inequality

$$\frac{1}{f\left(E^\theta(u_\theta(\mu,p^s))\right)}\le E^\theta\left(\frac{1}{f(u_\theta(\mu,p^s))}\right)=\frac{1}{f(\mu_\theta)}.$$

When $f(.)$ is increasing the extremes of this inequality imply that $\mu_\theta\le E^\theta(u_\theta(\mu,p^s))$. This establishes the first part of the result. Part (ii) is established by observing that the final inequality is reversed when concavity replaces convexity.

Finally, we must show that $\frac{1}{f(.)}$ cannot be concave on the open interval $(0,x)$, for any $x>0$. Suppose it were concave on such an interval for some $x\in(0,1)$ . Then for $\varepsilon<x$ and any $\lambda\in[0,1]$

$$\frac{1}{f(\lambda\varepsilon+(1-\lambda)x)}\ge\lambda\frac{1}{f(\varepsilon)}+(1-\lambda)\frac{1}{f(x)},$$

But as $\varepsilon\to0$ the RHS of the first of these inequalities converges to infinity (as $f(\varepsilon)\to0$). Thus $f((1-\lambda)x)=0$ for all $\lambda\in(0,1)$ which contradicts the fact that $f(.)$ is strictly increasing. $\square$

*Proof of Proposition 4*

*Proof.* A differentiable, strictly increasing map $f:[0,1]\to[0,1]$ determines the quasi-Bayesian updating $u_\theta(\mu,p^s)$ as in (16). We define the (differentiable and strictly increasing) function $\psi:\mathbb{R}\to\mathbb{R}$ as

$$\psi(\lambda)=\ln f\left(\frac{e^\lambda}{1+e^\lambda}\right)-\ln\left[1-f\left(\frac{e^\lambda}{1+e^\lambda}\right)\right],\qquad \lambda\in\mathbb{R}.$$

Also, note for later that

$$\psi'(\lambda)=f'(x)\frac{x}{f(x)}\frac{1-x}{1-f(x)},\qquad\text{where}\qquad x:=\frac{e^\lambda}{1+e^\lambda}. \tag{29}$$

If we define $\lambda_s':=\ln\frac{u_\theta(\mu,p_s)}{1-u_\theta(\mu,p_s)}$ and $\lambda:=\ln\frac{\mu_\theta}{1-\mu_\theta}$, then (16) implies that

$$\psi(\lambda_s')=\psi(\lambda)+\ln\frac{p_s^\theta}{p_s^{\theta'}}.$$

The function $\psi$ is invertible, so this allows us to write

$$\lambda_s'=\ln\frac{u_\theta(\mu,p_s)}{1-u_\theta(\mu,p_s)}=\psi^{-1}\left(\psi(\lambda)+\ln\frac{p_s^\theta}{p_s^{\theta'}}\right). \tag{30}$$

Now we calculate the variance of the updated beliefs log likelihood ratio. This is

$$\mathrm{Var}\left[\ln \frac{\mathrm{u}_\theta(\mu, \mathrm{p_s})}{1-\mathrm{u}_\theta(\mu, \mathrm{p_s})}\right] = \frac{1}{2}\sum_{\mathrm{s,s'\in S}} \pi_{\mathrm{s}}\pi_{\mathrm{s'}}\left(\ln \frac{\mathrm{u}_\theta(\mu, \mathrm{p_s})}{1-\mathrm{u}_\theta(\mu, \mathrm{p_s})} - \ln \frac{\mathrm{u}_\theta(\mu, \mathrm{p_{s'}})}{1-\mathrm{u}_\theta(\mu, \mathrm{p_{s'}})}\right)^2.$$

Here $\pi_s = \sum_\theta \mu^\theta p_s^\theta$ is defined to be the unconditional probability of the signal $s$ in the experiment $\mathcal{E}$. A substitution from (30) then gives

$$\mathrm{Var}\left[\ln \frac{\mathrm{u}_\theta(\mu, \mathrm{p_s})}{1-\mathrm{u}_\theta(\mu, \mathrm{p_s})}\right] = \frac{1}{2}\sum_{s,s'\in S} \pi_s\pi_{s'}\left(\psi^{-1}\left(\psi(\lambda) + \ln \frac{p_s^\theta}{p_s^{\theta'}}\right) - \psi^{-1}\left(\psi(\lambda) + \ln \frac{p_{s'}^\theta}{p_{s'}^{\theta'}}\right)\right)^2.$$

As we have assumed the function $f$ is continuously differentiable, we can apply the intermediate value theorem to the function $\psi^{-1}$. Hence,

$$\psi^{-1}\left(\psi(\lambda) + \ln \frac{p_s^\theta}{p_s^{\theta'}}\right) - \psi^{-1}\left(\psi(\lambda) + \ln \frac{p_{s'}^\theta}{p_{s'}^{\theta'}}\right) = \frac{d\psi^{-1}(\tilde{\lambda})}{d\lambda}\left(\ln \frac{p_s^\theta}{p_s^{\theta'}} - \ln \frac{p_{s'}^\theta}{p_{s'}^{\theta'}}\right)$$

for some $\tilde{\lambda}$ satisfying $\min_s \ln \frac{p_s^\theta}{p_s^{\theta'}} \leq \tilde{\lambda} - \psi(\lambda) \leq \max_s \ln \frac{p_s^\theta}{p_s^{\theta'}}$. Let $B$ denote this interval of potential values of $\tilde{\lambda}$. If this calculation is substituted into the expression for the variance we can then get a lower bound on the variance

$$\begin{aligned}\mathrm{Var}\left[\ln \frac{\mathrm{u}_\theta(\mu, \mathrm{p_s})}{1-\mathrm{u}_\theta(\mu, \mathrm{p_s})}\right] &\geq \min_{\tilde{\lambda}\in B}\left[\frac{d\psi^{-1}(\tilde{\lambda})}{d\lambda}\right]^2 \frac{1}{2}\sum_{s,s'\in S}\pi_s\pi_{s'}\left(\ln \frac{p_s^\theta}{p_s^{\theta'}} - \ln \frac{p_{s'}^\theta}{p_{s'}^{\theta'}}\right)^2 \\ &= \min_{\tilde{\lambda}\in B}\left[\frac{d\psi^{-1}(\tilde{\lambda})}{d\lambda}\right]^2 \mathrm{Var}\left[\ln \frac{\mathrm{p}^\theta}{\mathrm{p}^{\theta'}}\right].\end{aligned}$$

An upper bound can be obtained in a similar way

$$\mathrm{Var}\left[\ln \frac{\mathrm{u}_\theta(\mu, \mathrm{p_s})}{1-\mathrm{u}_\theta(\mu, \mathrm{p_s})}\right] \leq \max_{\tilde{\lambda}\in B}\left[\frac{d\psi^{-1}(\tilde{\lambda})}{d\lambda}\right]^2 \mathrm{Var}\left[\ln \frac{\mathrm{p}^\theta}{\mathrm{p}^{\theta'}}\right].$$

These inequalities imply that bounding the derivatives of $\psi^{-1}$ will generate over and under reaction. As $\psi$ and its inverse are strictly increasing functions with positive derivatives. The above inequalities imply that a sufficient condition for the updating to satisfy the condition for overreaction is $d\psi^{-1}(\tilde{\lambda})/d\lambda > 1$ for all $\tilde{\lambda}$ and a sufficient condition for under-reaction is $d\psi^{-1}(\tilde{\lambda})/d\lambda < 1$. The calculation of the derivative $d\psi/d\lambda$ in (29) then implies the sufficient conditions given in the Proposition. □

*Proof of Proposition 5*

*Proof.* We begin by redescribing the updating using matrix notation: The experiment $\mathcal{E}_n$ is denoted by a non-negative $n\times|\Theta|$ matrix $P$ with columns $p^\theta \in \Delta(S_n)$ (the signal distributions for the states $\theta$). The prior $\mu \in \Delta(\Theta)$ is a column vector. The vector of (unconditional) signal probabilities in the experiment equals $P\mu \in \Delta(S_n)$. The updating function, $\mathcal{U}_n(\mu, \mathcal{E})$, is denoted by the $|\Theta| \times n$ matrix $U = \mathcal{U}(\mu, P)$, with columns $(\hat{\mu}_s)_{s\in S_n}$ that are the updated beliefs after the signal $s$. The updating function is unbiased iff:

$$\mu \equiv \mathcal{U}(\mu, P)P\mu; \qquad \text{for all} \qquad \mu \in \Delta(\Theta), P \in \Delta(\Theta)^n. \tag{31}$$

Note that $\mathbb{1}^T\mathcal{U}(\mu,P)=\mathbb{1}^T$ for all $\mu$, $P$.

We now describe Bayesian updating using this notation. Let $D_{(x)}$ denote the $m\times m$ matrix with the vector $x\in\mathbb{R}^m$ on its diagonal and zeros elsewhere. Given the experiment $(\mu,P)$, the $|\Theta|\times n$ matrix $D_{(\mu)}P^T$ is the joint distribution of signals and states. (Its $\theta^{\text{th}}$ row is the unconditional probability of signal $s$ and state $\theta$.) As the unconditional probability of the signals is $P\mu$, the Bayesian update given prior $\mu$ and experiment $P$ is given by the function

$$\mathcal{U}^B(\mu,P):=D_{(\mu)}P^TD^{-1}_{(P\mu)}. \tag{32}$$

We now can begin the proof with the "only if" direction. Let us suppose that $\mathcal{U}(\mu,P)=\mathcal{U}^B(\mu,Q)$ for some $Q\in\Delta(S_n)^{|\Theta|}$ satisfying $P\mu=Q\mu$. To show that $\mathcal{U}$ is unbiased, that is it satisfies (31), a substitution from (32) gives

$$\mathcal{U}(\mu,P)P\mu=\mathcal{U}^B(\mu,Q)P\mu=D_{(\mu)}Q^TD^{-1}_{(Q\mu)}P\mu.$$

As $P\mu=Q\mu$ this then can be written as

$$\mathcal{U}(\mu,P)P\mu=D_{(\mu)}Q^TD^{-1}_{(Q\mu)}Q\mu=D_{(\mu)}Q^T\mathbb{1}=D_{(\mu)}\mathbb{1}=\mu.$$

Hence we have established $\mathcal{U}(\mu,P)P\mu=\mu$ and the updating is unbiased.

We now do the "if" direction. Suppose that $\mathcal{U}(\mu,P)$ is unbiased and satisfies (31). We will derive from it a new experiment $Q$ with the same unconditional signal probabilities.

$$\begin{aligned}
\mu&=\mathcal{U}(\mu,P)P\mu\\
D^{-1}_{(\mu)}\mu&=D^{-1}_{(\mu)}\mathcal{U}(\mu,P)P\mu\\
\mathbb{1}&=D^{-1}_{(\mu)}\mathcal{U}(\mu,P)D_{(P\mu)}D^{-1}_{(P\mu)}P\mu\\
\mathbb{1}&=\underbrace{D^{-1}_{(\mu)}\mathcal{U}(\mu,P)D_{(P\mu)}}_{=:Q^T}\mathbb{1}
\end{aligned}$$

The final equality above says that $Q$ is a $n\times|\Theta|$ matrix with columns that are probabilities on $S$. Thus it is a feasible experiment. The unconditional signal probabilities in $Q$ are the same as $P$, because

$$Q\mu=D_{(P\mu)}\mathcal{U}(\mu,P)^TD^{-1}_{(\mu)}\mu=D_{(P\mu)}\mathcal{U}(\mu,P)^T\mathbb{1}=D_{(P\mu)}\mathbb{1}=P\mu.$$

Furthermore, the Bayesian update with the experiment $Q$ and initial beliefs $\mu$ equals $\mathcal{U}(\mu,P)$ because

$$\mathcal{U}^B(\mu,Q)=D_{(\mu)}Q^TD^{-1}_{(Q\mu)}=D_{(\mu)}[D^{-1}_{(\mu)}\mathcal{U}(\mu,P)D_{(P\mu)}]D^{-1}_{(P\mu)}=\mathcal{U}(\mu,P).$$

(The second equality substitutes for $Q^T$ and $P\mu=Q\mu$.) Thus Bayesian updating with the experiment $Q$ gives the same update as the arbitrary updating with experiment $P$. □

*Proof of Proposition 6*

*Proof.* The "if" part of the proof is trivial. It is well known that Bayesian updating satisfies all the claimed properties.

The "only if" part proceeds in several stages. In Part 1 Axiom 6 is applied to a particular binary experiment. To show that the homeomorphism maps rays from $e_\theta$ to other rays from from $e_\theta$. Then (in Part 2), we show that $F_{\theta'}(\mu)/F_{\theta''}(\mu) \equiv \psi_{\theta'}(\mu_\theta)/\psi_{\theta''}(\mu_{\theta''})$ and these ratios are homogeneous degree zero. This holds as the ratio $F_{\theta'}(\mu)/F_{\theta''}(\mu)$ is constant along these rays. This allows us to show (in Part 3) that $F$ is the homeomorphism that generates geometric probability weighting. Finally, (Part 4) we use a linear dependence argument to show that the only geometric probability weighting updating that satisfies Axiom 6 is Bayesian updating.

Part 1. By Proposition 1, the updating is described by (4) and a homeomorphism $F$. We begin by deriving the property (33) that $F$ satisfies if Axiom 6 holds. Let $e_\theta$ denote the vector with unity in the $\theta$ element and zeros elsewhere. When the agent observes a binary experiment with the vectors of signal probabilities $p_1 = ze_\theta$ and $p_2 = \mathbb{1} - ze_\theta$, $z \in (0,1)$, Axiom 6 implies that

$$\mu \in \text{Convex Hull}\left\{F^{-1}(e_\theta), F^{-1}\left(\frac{F(\mu)\circ(\mathbb{1} - ze_\theta)}{F(\mu)^T(\mathbb{1} - ze_\theta)}\right)\right\}.$$

But, $F(e_\theta) = e_\theta$ by (15) and a substitution into (4). So, for all $z \in (0,1)$ and $\mu \neq e_\theta$ there exists an $\omega \in [0,1)$ such that

$$F^{-1}\left(\frac{F(\mu)\circ(\mathbb{1} - ze_\theta)}{F(\mu)^T(\mathbb{1} - ze_\theta)}\right) = \frac{\mu - \omega e_\theta}{1-\omega}.$$

Now apply $F$ to both sides of the above and substitute $F(\mu)^T(\mathbb{1} - ze_\theta) = 1 - zF_\theta(\mu)$. This gives the equivalent condition that for all $z \in (0,1)$ there exists an $\omega \in [0,1)$ such that

$$\frac{F(\mu)\circ(\mathbb{1} - ze_\theta)}{1 - zF_\theta(\mu)} = F\left(\frac{\mu - \omega e_\theta}{1-\omega}\right) := F(\tilde{\mu}). \tag{33}$$

(33) says that $F$ maps line segments in $\Delta(\Theta)$ to line segments in $\Delta(\Theta)$. The point $\tilde{\mu}$ lies on the line segment joining $\mu$ to $(0, \frac{\mu_{-\theta}}{\mathbb{1}^T\mu_{-\theta}})$. By increasing $\mu_\theta$ while holding $\frac{\mu_{-\theta}}{\mathbb{1}^T\mu_{-\theta}}$ constant, $\tilde{\mu}$ can be any convex combination of $e_\theta$ and $(0, \frac{\mu_{-\theta}}{\mathbb{1}^T\mu_{-\theta}})$. This is a ray from $e_\theta$. This ray is mapped by $F$ to the ray with the extreme points $e_\theta = F(e_\theta)$ (when $\mu \to e_\theta$) and $\frac{(0, F_{-\theta}(\mu))}{1-F_\theta(\mu)}$ $(z \to 1)$,[36] as

$$\frac{F(\mu)\circ(\mathbb{1} - ze_\theta)}{1 - zF_\theta(\mu)} = \delta e_\theta + (1-\delta)\frac{(0, F_{-\theta}(\mu))}{1 - F_\theta(\mu)},$$

where $\delta = (1-z)F_\theta(\mu)(1 - zF_\theta(\mu))^{-1} \in [0,1]$. $F$ is a homeomorphism, so it is also a homeomorphism when it is restricted to these rays. Thus $F$ maps $e_\theta$ to $e_\theta$ and $(0, \frac{\mu_{-\theta}}{\mathbb{1}^T\mu_{-\theta}})$ to $\frac{(0, F_{-\theta}(\mu))}{1-F_\theta(\mu)}$. Furthermore, as $(0, \frac{\mu_{-\theta}}{\mathbb{1}^T\mu_{-\theta}}) \in \Delta(\Theta)$ varies these rays range over all $\tilde{\mu}$ in $\Delta(\Theta)$.

Part 2. We now show that $F_{\theta'}(\mu)/F_{\theta''}(\mu) \equiv A_{\theta',\theta''}(\mu_{\theta'}, \mu_{\theta''})$ where $A_{\theta',\theta''}(.)$ is homogeneous degree zero. That is, this ratio only depends on the variables $\mu_{\theta'}$ and $\mu_{\theta''}$. For any $\theta', \theta'' \neq \theta$ a substitution from (33) gives

$$\frac{F_{\theta'}(\tilde{\mu})}{F_{\theta''}(\tilde{\mu})} = \frac{\frac{F_{\theta'}(\mu)}{1-zF_\theta(\mu)}}{\frac{F_{\theta''}(\mu)}{1-zF_\theta(\mu)}} = \frac{F_{\theta'}(\mu)}{F_{\theta''}(\mu)}, \qquad \forall z, \omega.$$

[36] $F_{-\theta}(\mu)$ denotes $F(\mu)$ with its $\theta$ entry missing.

Thus the ratio $F_{\theta'}(\tilde{\mu})/F_{\theta''}(\tilde{\mu})$ is constant for all $\tilde{\mu} = (\mu - \omega e_\theta)(1-\omega)^{-1}$ and for all $\theta \neq \theta', \theta''$. This is because the RHS of the above is independent of $\omega$ or $z$. Let $\hat{\mu}, \breve{\mu} \in \Delta^o(\Theta)$ have $\theta'$ and $\theta''$ entries that are collinear: $(\hat{\mu}_{\theta'}, \hat{\mu}_{\theta''}) = \kappa(\breve{\mu}_{\theta'}, \breve{\mu}_{\theta''})$. We need to show that

$$\frac{F_{\theta'}(\hat{\mu})}{F_{\theta''}(\hat{\mu})} = \frac{F_{\theta'}(\breve{\mu})}{F_{\theta''}(\breve{\mu})}. \tag{34}$$

We choose $\bar{\mu} \in \Delta^o(\Theta)$ such that $(\bar{\mu}_{\theta'}, \bar{\mu}_{\theta''}) = \bar{\kappa}(\breve{\mu}_{\theta'}, \breve{\mu}_{\theta''})$ and $\bar{\mu}_\theta/\bar{\mu}_{\theta'} > \max\{\tilde{\mu}_\theta/\tilde{\mu}_{\theta'}, \breve{\mu}_\theta/\breve{\mu}_{\theta'}\}$ for all $\theta \neq \theta', \theta''$. This is possible as $\bar{\kappa} \geq \bar{\mu}_{\theta'}$ can be chosen to be arbitrarily small. Now notice that:

$$\hat{\mu} = \frac{\bar{\mu} - \sum_{\theta \neq \theta', \theta''} \hat{\omega}_\theta e_\theta}{1 - \sum_{\theta \neq \theta', \theta''} \hat{\omega}_\theta}, \qquad \breve{\mu} = \frac{\bar{\mu} - \sum_{\theta \neq \theta', \theta''} \breve{\omega}_\theta e_\theta}{1 - \sum_{\theta \neq \theta', \theta''} \breve{\omega}_\theta};$$

where $\hat{\omega}_\theta := \bar{\mu}_\theta - \hat{\mu}_\theta(\bar{\mu}_{\theta'}/\hat{\mu}_{\theta'})$ and $\breve{\omega}_\theta := \bar{\mu}_\theta - \breve{\mu}_\theta(\bar{\mu}_{\theta'}/\breve{\mu}_{\theta'})$. The assumptions on $\bar{\mu}$ ensure that $\hat{\omega}_\theta, \breve{\omega}_\theta > 0$ for all $\theta \neq \theta', \theta''$ and that the denominators above are positive. Thus both of $\hat{\mu}$ and $\breve{\mu}$ can be derived from $\bar{\mu}$ by a sequence of linear transformations of the form $(\mu - \omega e_\theta)/(1-\omega)$; one for each $\theta \neq \theta', \theta''$. These transformations do not alter the ratio $F_{\theta'}(\mu)/F_{\theta''}(\mu)$ hence (34) is established.

Now it is shown that $A_{\theta',\theta''}(\mu_{\theta'}, \mu_{\theta''}) = \psi_{\theta'}(\mu_{\theta'})/\psi_{\theta''}(\mu_{\theta''})$ for all $\theta', \theta'' \in \Theta$. The proposition asserts that there exists at least three different values $\theta, \theta', \theta'' \in \Theta$. Thus we can find functions $A_{\theta,\theta'}$, $A_{\theta',\theta'}$, and $A_{\theta,\theta''}$ such that

$$\frac{F_\theta(\mu)}{F_{\theta'}(\mu)} \frac{F_{\theta'}(\mu)}{F_{\theta''}(\mu)} \equiv \frac{F_\theta(\mu)}{F_{\theta''}(\mu)}, \qquad \text{or} \qquad A_{\theta,\theta'}(\mu_\theta, \mu_{\theta'}) A_{\theta',\theta''}(\mu_{\theta'}, \mu_{\theta''}) \equiv A_{\theta,\theta''}(\mu_\theta, \mu_{\theta''}).$$

Dividing by $A_{\theta',\theta''}(\mu_{\theta'}, \mu_{\theta''})$ gives

$$A_{\theta,\theta'}(\mu_\theta, \mu_{\theta'}) \equiv \frac{A_{\theta,\theta''}(\mu_\theta, \mu_{\theta''})}{A_{\theta',\theta''}(\mu_{\theta'}, \mu_{\theta''})}.$$

The left above does not depend on $\mu_{\theta''}$ so neither does the right. Fixing the value of $\mu_{\theta''} = t$ and defining $\psi_\theta(\mu_\theta) := A_{\theta,\theta''}(\mu_\theta, t)$ and $\psi_{\theta'}(\mu_{\theta'}) := A_{\theta',\theta''}(\mu_{\theta'}, t)$ we get that

$$A_{\theta,\theta'}(\mu_\theta, \mu_{\theta'}) \equiv \frac{\psi_\theta(\mu_\theta)}{\psi_{\theta'}(\mu_{\theta'})}.$$

As $\theta$ and $\theta'$ are entirely arbitrary we have established the claim, that is,

$$\frac{F_\theta(\mu)}{F_{\theta'}(\mu)} = A_{\theta,\theta'}(\mu_\theta, \mu_{\theta'}) = \frac{\psi_\theta(\mu_\theta)}{\psi_{\theta'}(\mu_{\theta'})} \qquad \forall \theta, \theta' \in \Theta. \tag{35}$$

Part 3. We now derive and solve functional (Pexider) equations that the $\psi_\theta$ satisfy. This will tell us the form of $\psi_\theta$. The homogeneity degree zero of (34) says that

$$\frac{\psi_\theta(\mu_\theta)}{\psi_{\theta'}(\mu_{\theta'})} = \frac{\psi_\theta(\rho\mu_\theta)}{\psi_{\theta'}(\rho\mu_{\theta'})}, \qquad \text{or} \qquad \frac{\psi_{\theta'}(\rho\mu_{\theta'})}{\psi_{\theta'}(\mu_{\theta'})} = \frac{\psi_\theta(\rho\mu_\theta)}{\psi_\theta(\mu_\theta)}.$$

for all $\rho \in (0,1)$. If we define $g_\theta := \ln \psi_\theta$, $m_\theta := \ln \mu_\theta$, and $\ell := \ln \rho$, then this can be rewritten as

$$g_{\theta'}(e^{\ell + m_{\theta'}}) - g_{\theta'}(e^{m_{\theta'}}) = g_\theta(e^{\ell + m_\theta}) - g_\theta(e^{m_\theta}).$$

Another change of variable, letting $f_\theta(x) := g_\theta(e^x)$, and another rewriting gives

$$f_{\theta'}(\ell + m_{\theta'}) - f_{\theta'}(m_{\theta'}) = f_\theta(\ell + m_\theta) - f_\theta(m_\theta).$$

This Pexider equation, in $f_\theta$ and $f_{\theta'}$, can be solved in the following way. Choose $m_\theta = m_{\theta'}$ and $x = \ell + m_\theta$ then $f(x) := f_\theta(x) = f_{\theta'}(x) + c_{\theta'}$ where $c_{\theta'}$ is a constant. Rewriting the above using $f$, setting $m_{\theta'} = -\ell$ and re-arranging a little gives

$$f(m_\theta) - f(0) = f(\ell + m_\theta) - f(0) + f(-\ell) - f(0).$$

Thus the function $f(m_\theta) - f(0)$ satisfies Cauchy's functional equation and the only continuous solution to this is $f(m_\theta) - f(0) = a m_\theta$ for some constant $a$.[37] This implies that $f_\theta(m_\theta) = a m_\theta + c_\theta$ for all $\theta$. Then, reversing all the previous transformations gives

$$\psi_\theta(\mu_\theta) = K_\theta \mu_\theta^a;$$

for some $K_\theta$ and all $\theta$. Finally, as $\psi_\theta/\psi_{\theta'} = F_\theta/F_{\theta'}$ and $\mathbb{1}^T F(\mu) = 1$ we get

$$F(\mu) = \frac{\left(K_1 \mu_1^a, \ldots, K_{|\Theta|}\mu_{|\Theta|}^a\right)}{\sum_\theta K_\theta \mu_\theta^a}. \tag{36}$$

This homeomorphism generates the geometric probability weighting updating with the weight $a$. The formula for the updating $F$ generates is given in Section 4.1.

Part 4. The final step in the proof is to show that $a = 1$. Suppose that: $F$ has the form (36), the agent has the initial belief $\mu$, and she observes the binary experiment with signal probabilities $p_1 \in (0,1)^{|\Theta|}$ and $\mathbb{1} - p_1$. If we apply the updating formula of Section 4.1 to Axiom 6 in this case we get the condition

$$\mu \in \text{Convex Hull}\left\{\frac{\mu \circ (p_1)^{\frac{1}{a}}}{\mu^T (p_1)^{\frac{1}{a}}}, \frac{\mu \circ (\mathbb{1} - p_1)^{\frac{1}{a}}}{\mu^T (\mathbb{1} - p_1)^{\frac{1}{a}}}\right\},$$

where $(p_1)^{\frac{1}{a}} = ((p_1^1)^{1/a}, \ldots, (p_1^{|\Theta|})^{1/a})$. This implies that there exists $w_1, w_2 \in \mathbb{R}$ such that

$$0 = \mu - w_1 \frac{\mu \circ (p_1)^{\frac{1}{a}}}{\mu^T (p_1)^{\frac{1}{a}}} - w_2 \frac{\mu \circ (\mathbb{1} - p_1)^{\frac{1}{a}}}{\mu^T (\mathbb{1} - p_1)^{\frac{1}{a}}}$$

or that

$$0 = \mu \circ \left(\mathbb{1} - (p_1)^{\frac{1}{a}} \frac{w_1}{\mu^T (p_1)^{\frac{1}{a}}} - (\mathbb{1} - p_1)^{\frac{1}{a}} \frac{w_2}{\mu^T (\mathbb{1} - p_1)^{\frac{1}{a}}}\right).$$

Thus the vectors $\mathbb{1}$, $(p_1)^{\frac{1}{a}}$, and $(\mathbb{1} - p_1)^{\frac{1}{a}}$ are linearly dependent for any $p_1 \in (0,1)^{|\Theta|}$. If $|\Theta| > 2$, this implies $a = 1$. If $a = 1$ the agent is updating as a Bayesian and the result is established. $\square$

[37] The updating is assumed to be continuous.