

DIVISIBLE UPDATING

MARTIN W. CRIPPS

ABSTRACT. A characterisation is provided of the belief updating processes that are independent of how an individual chooses to divide up/partition the statistical information they use in their updating. These “divisible” updating processes are in general not Bayesian, but can be interpreted as a re-parameterisation of Bayesian updating. This class of rules incorporates over- and under-reaction to new information in the updating and other biases. We also show that a martingale property is, then, sufficient the updating process to be Bayesian.

1. BELIEF UPDATING

In this paper we consider arbitrary processes for updating beliefs in the light of statistical evidence. We treat these updating processes as deterministic maps from the prior beliefs and a signal structure (statistical experiments) to a profile of updated beliefs (one for each possible signal). We place axioms on these updating processes and show that a generalisation of Bayesian updating can be derived as a consequence of these axioms.

Consider an individual who receives two signals/pieces of information/news. There are several ways such an individual can use these two pieces of information to update their beliefs about the world. One is to consider the joint distribution of these pieces of information and to do just one update. An alternative (which is natural when information arrives sequentially or over time) is to separate the two pieces of information and to update beliefs twice. That is, to update beliefs once using the first piece of information and its distribution. And then to update these intermediate beliefs a second time using the second piece of information and its conditional distribution given the first piece of information. If these two different procedures generate the same ultimate profile of updated beliefs we will say that the updating is divisible. We will characterise all the updating processes that have this divisibility property and show that they can be interpreted as natural generalisation of Bayesian updating. Furthermore we will show that divisibility plus an unbiasedness/martingale property for the updating uniquely characterises Bayesian updating.

The main result of this paper shows that if updating satisfies four properties, then it is characterised by a map F from the space of beliefs into itself. The four properties or axioms are: First, that uninformative experiments do not result in changes in beliefs.

Date: March 15, 2018.

Department of Economics, University College London; m.cripps@ucl.ac.uk . My thanks are due to Aislinn Bohren, Toru Kitagawa, and Ran Speigler for their suggestions and comments.

Second, that the names of the signals do not matter just their probability content. The third ensures that all updated beliefs are possible. The fourth is the divisibility property described above. Any updating procedure, that satisfies these four properties, follows 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 map F^{-1} .



FIGURE 1. Updating that Satisfies Divisibility.

Our result is that if the belief revision protocol satisfies the divisibility property, then it must have this structure. (Where F is a homeomorphism from the space of beliefs into itself.) It is clear that Bayesian updating is a subset of this class, (F is the identity). Specific properties of F ensure different properties of the belief revision. Biases in the updating will be generated by the convexity or concavity of a function of the map F . For overreaction and under-reaction to signals we will argue that F needs to be an expansive or a contraction mapping. In regions where F is a contraction the belief revision will exhibit over-reaction in general and in regions where F^{-1} is a contraction the belief revision will exhibit underreaction.

Many of the useful properties of Bayesian updating also carry over to this class of revision protocols. For example, although the actual beliefs are not a martingale the shadow beliefs are. Thus, in dynamic settings, updating that satisfies divisibility will be a deterministic function of a martingale. As a result, consistency will hold for this larger class of updating processes, provided F maps the extreme points of the belief simplex to themselves. That is, when these updating processes are repeatedly applied to a sequence of data, they will generate limiting beliefs that attach probability one to the truth. This property of divisible updating contrasts with other models of non-Bayesian updating, for example Rabin and Schrag (1999) and Epstein, Noor, and Sandroni (2010).

A question we also address is what additional axioms are necessary to characterise Bayesian updating. We show in the final section of this paper that only one further property is sufficient for this, that is, belief revision is unbiased or follow a martingale.¹ Thus any relaxation of Bayesian updating has to either violate the martingale property, the divisibility property or one of the other axioms we impose.

Why do we focus on the divisibility of updating rather than some other property it may have? There are several motivations for looking at divisibility. First, the information individuals observe is typically not a unitary piece of information but a bundle of signals. What divisibility ensures is that the exact way the updating is performed (using this

¹There are several names used for this notion in different contexts: unbiased, Bayes plausibility have also been used.

bundle of signals) is unimportant in determining the final beliefs. And, if the individual separates out various features of the bundle and uses these individually to update their beliefs, then the order that they do this has no ultimate effect their eventual beliefs. If updating does not satisfy the divisibility property, then the updated beliefs will depend on how and in what order the individual has applied the updating (see Section 1.1 below for an example of this). Moreover, if divisibility is not satisfied, the individual has multiple possible updated beliefs and an additional assumption is required to pin them down. To predict the individual's updated beliefs it will then also be necessary to have a theory of how she chose to apply the updating rules to packages of signals. Furthermore, she may also benefit by recalling more of the history than her current state of beliefs if the order of the updating matters.

In dynamic settings this issue is most clear. If the updating satisfies divisibility, then it is only necessary for the individual to keep track of their current beliefs and update them any time new information appears. Her current state of beliefs are sufficient to summarise all past signals in a parsimonious way—nothing else about the past history must be remembered. There is a dynamic consistency in the updating so the individuals who are not required to act can simply collect information and use it to update beliefs when it is necessary. If the updating is not divisible, then how often the updating rule has been applied in the past may affect her current beliefs and thus may be something she needs to keep track of. In summary, if the belief updating is not divisible, then a theory of how and when updating occurs is required. Such a theory needs to address the trade-offs between the memory costs of storing accumulated signals and the processing costs of when and how to update.

1.1. *An Example of Non-Divisible Updating*

In this section we consider the model of non-Bayesian learning described Epstein, Noor, and Sandroni (2010) or conservative Bayesian updating of Hagmann and Loewenstein (2017) and apply it to an example.² In these models of learning, the updated beliefs are a linear combination of the prior and the posterior. We will show, using an example, that this model of updating does not satisfy the divisibility property and seek to explain why the updated beliefs vary for different protocols.

Consider an individual who is waiting for a bus and is learning about 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 $\mu \in (0, 1)$ that the state is good and a bus will arrive at some point. If no bus arrives in period $t = 0$, then a Bayesian would revise these beliefs downward in the light of this bad news to $\frac{\alpha\mu}{(1-\mu)\beta + \mu\alpha}$. However, in this model of non-Bayesian learning

²Both of these papers use this model of updating as a part of more general study of dynamic choice when there is learning or information design. We do not address choice issues here, focussing solely on belief revision.

she revises her beliefs to a weighted average of the prior and Bayesian posterior:

$$(1) \quad \mu_1 = (1 - \lambda)\mu + \lambda \frac{\alpha\mu}{(1 - \mu)\beta + \mu\alpha}, \quad \lambda \geq 0.$$

This is a particularly elegant model as it preserves the martingale property of the Bayesian posteriors. The value λ is a parameter of the updating procedure, which can in general depend on the current history, but in this example it is treated as a fixed property of the updating. This model of updating generalises Bayesian updating ($\lambda = 1$), where λ adjusts the effect the bad news of no arrival has on the beliefs. Choosing $\lambda < 1$ allows the individual to be under-confident or conservative in their updating of their beliefs. Conversely, $\lambda > 1$ allows the individual to be overconfident about the new information they have received. Thus, λ allows a range of updates from extreme dogmatism ($\lambda = 0$) or jumping to certainty ($\lambda \rightarrow \mu^{-1}(1 - \alpha)^{-1}$).

Now consider what the individual's beliefs could be after t periods waiting for a bus without an arrival. One possible application of this updating procedure is to iterate the updating process described in (1) for each of the t periods the individual has been waiting, to get

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

This updated belief in period t , μ_t , clearly has required considerable mental agility on the part of the individual. But it requires less memory, as the individual does not need to keep track of all past events—the current value of the beliefs μ_s is all that she needs to know.

An alternative application of the updating procedure (1) would be to suppose that the individual took all their current information at time t and performed one update of their prior. This is the update (for example) that is done by someone who only gets to the bus stop in period t and learns a bus has not yet arrived. Such a person would do one Bayesian update and arrive at the updated belief $\frac{\alpha^t\mu}{(1-\mu)\beta^t + \mu\alpha^t}$.³ Hence the non-Bayesian updating protocol (1) would in this case have an updated belief:

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

The updated belief $\tilde{\mu}_t$ has required a different kind of mental agility—requiring the individual to keep track of the past history of events and their original prior. The feature of the updating procedure (1) we seek to remedy below is that the iterated and one-shot update result in different beliefs: $\tilde{\mu}_t \neq \mu_t$. The two different ways of using the same information has generated different updated beliefs. The only time $\tilde{\mu}_t = \mu_t$ is when $\lambda = 1$ and the updating is fully Bayesian.

There is no simple comparative static that describes when one of the updating procedures described above makes the individual more optimistic than the other. If we take the initial beliefs $\mu = \frac{1}{2}$ for simplicity, then $\tilde{\mu}_2 = (1 - \lambda)\frac{1}{2} + \lambda \frac{\alpha^2}{\beta^2 + \alpha^2}$ and the iterated updating gives: $\mu_1 = (1 - \lambda)\frac{1}{2} + \lambda \frac{\alpha}{\beta + \alpha}$, and $\mu_2 = (1 - \lambda)\mu_1 + \lambda \frac{\alpha\mu_1}{(1 - \mu_1)\beta + \mu_1\alpha}$. A calculation

³ α^t is the probability that no bus arrives in the first t periods in the good state.

(see Appendix) shows that

$$\frac{\tilde{\mu}_2 - \mu_2}{\lambda\gamma(1-\gamma)} = \frac{1}{\frac{1}{\mu_1} + \frac{\gamma}{1-\mu_1}} - \frac{1}{1 + \gamma + \gamma^2 + \gamma^3}, \quad \gamma := \frac{\alpha}{\beta}, \quad \mu_1 \in (0, 1/2).$$

This implies that $\tilde{\mu}_2 > \mu_2$ if and only if $\lambda < 1$. The intuition for this comparative static is as follows. Varying λ tends to emphasise either the data or the prior. The iteration of the updating mitigates this effect. To see how this works in practice first suppose that $\lambda < 1$ and the prior is given increased weight. In the one-shot update, $\tilde{\mu}_2$, there is only one opportunity for the bad news—two periods without busses—to drive the prior down. But, when updating is iterated, it weakens the effect of the increased weight on the original prior. It, instead, places some weight on the intermediate prior, μ_1 . Hence, μ_2 will tend to be smaller than $\tilde{\mu}_2$ when the information that arrives is bad news, because μ_2 places less weight on the original prior than $\tilde{\mu}_2$.

Conversely, when $\lambda > 1$ the data is given increased weight in the updating. The one-shot update, $\tilde{\mu}_2$, gives these two periods of bad news excessive weight. Whereas the iterated update, μ_2 , decreases the weight given to the first period of the bad news by averaging it with the prior. So, the individual is less optimistic when they do a one-shot up date because it maximises the weight given to the data that has been collected. The iteration now dilutes the effect of the over-emphasized data.⁴

Of course, the two ways of determining period t beliefs, μ_t and $\tilde{\mu}_t$, are not the only ways the beliefs in period t could be arrived at. An individual could update several times over the periods $s = 0, \dots, t - 1$ —how many times this occurred would depend on the agent's costs of cogitation. Thus there is a whole family of potential updated beliefs at time t and as time passes this family grows.

1.2. Related Literature

There is a growing Economics literature, both experimental and theoretic, investigating the consequences of a non-Bayesian updating of beliefs, see for example: Rabin and Schrag (1999), Ortoleva (2012), Angrisani, Guarino, Jehiel, and Kitagawa (2017), Levy and Razin (2017), Brunnermeier (2009), Bohren and Hauser (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 properties of a particular assumption about how updating may occur. 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 has already been considered at length in the example above.

⁴When t varies continuously it is possible to get a cleaner expression for μ_t and a similar comparative static on μ_t and $\tilde{\mu}_t$ can be obtained. This is also done in the Appendix.

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), has perhaps the closest connection with this one. Here it is shown 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, a property equivalent to divisibility. These reweighted Bayesian updates are a subset of the class of updating rules that described here (the rules that update as in Figure 1). Here we described all continuous updating rules that satisfy commutativity/divisibility. In Zhao (2016) a set of weaker axioms are shown to characterise an updating rule that does not satisfy divisibility, but does satisfy an order independence property similar to the one discussed in the appendix, however, this property is required to hold only for independent events.

There are links between the notion of divisibility and the models dynamic choice under uncertainty. In particular the literature on recursive preferences in dynamic settings, or dynamically consistent preference update rules. Here agents are required to act consistently in situations where information arrives over many periods and thus implicitly behave as if they update divisibly; see for example Epstein and Zin (1989) or Epstein and Schneider (2003)).

In the statistics literature the notion of prequentiality, Dawid (1984), emphasises the idea that that forecasting should be an iterative procedure and there may be differences between iteratively revised forecasts and other kinds of forecasting procedures.

2. A MODEL OF BELIEF UPDATING

In this section we describe our model of belief revision and the most important axioms we will impose on the updating. The approach taken is inspired by the axiomatic interpretation of entropies: see for example Shannon and Weaver (1949), Tverberg (1958) or Aczél and Dhombres (1989) p.66. We will not adopt the terminology of “priors” and “posteriors”, reserving these terms for the Bayesian updating only. Instead the agent is assumed to be equipped with “beliefs” that are revised when information arrives to form “updated beliefs”.

There are K possible values for an unknown parameter $\theta \in \{1, 2, \dots, K\} := \Theta$ and the agent has the initial beliefs, $\mu = (\mu^1, \dots, \mu^K) \in \Delta(\Theta)$, about the value of this parameter. There is a statistical experiment \mathcal{E} that the agent can conduct which provides further information on θ .⁵ The experiment comprises a finite set of signals $s \in \{1, 2, \dots, n\} = S$ and parameter-dependent probabilities for the signals $p^\theta = (p_1^\theta, \dots, p_n^\theta)$. We will only consider experiments with strictly positive probabilities for all signals. Hence, we assume $p_s^\theta > 0$ for all $s \in S$ and $\theta \in \Theta$, or that $p^\theta \in \Delta^\circ(S)$.⁶ We will also want to consider the probabilities of a given signal $s \in S$, hence we will define $p_s := (p_s^1, \dots, p_s^K) \in (0, 1)^K$ to be these probabilities. Thus p^θ are the rows of the matrix \mathcal{E} and p_s are its columns. In

⁵See Torgersen (1991), for example, for the general properties of statistical experiments.

⁶We use $\Delta^\circ(S)$ to denote the interior of $\Delta(S)$.

summary, the agent has priors $\mu \in \Delta(\Theta)$ and access to the experiment $\mathcal{E} := ((p^\theta)_{\theta \in \Theta}) \in \Delta^o(S)^K$.

An updating process takes the experiment, \mathcal{E} , and the beliefs, μ , and maps them to the updated beliefs for each possible signal outcome, $s \in \{1, 2, \dots, n\}$. The outcome of the updating process is a profile of n possible updated beliefs $\{\mu^s \in \Delta(\Theta) : s \in S\}$ (one for each possible signal realisation $s \in S$). We define the updating function U_n to be the map from the beliefs and the experiment to the profile of updated beliefs, that is

$$U_n : \Delta(\Theta) \times \Delta^o(S)^K \rightarrow \Delta(\Theta)^n, \quad \text{for } n = 2, 3, \dots$$

We will also write $(\mu^1, \dots, \mu^n) = U_n(\mu, (p^\theta)_{\theta \in \Theta}) = U_n(\mu, \mathcal{E})$.⁷

We will impose some conditions that the updating must satisfy. The first condition is that if the signals are identical then there is no updating. That is if $p^\theta = p^{\theta'}$ for all $\theta, \theta' \in \Theta$ then the updated belief is the same as the original beliefs.

Axiom 1 (Uninformativeness). $U_n(\mu, \mathcal{E}) = (\mu, \dots, \mu)$, if $p^\theta = p^{\theta'}$ for all $\theta, \theta' \in \Theta$.

The second condition on the updating is that the names of the signals are unimportant for how the beliefs are revised, it is only the probabilities in the experiment that matter. Thus permuting the order of the signals just permutes the order of the updated probabilities.

Axiom 2 (Symmetry). For any n , any permutation $\omega : \{1, 2, \dots, n\} \rightarrow \{1, 2, \dots, n\}$, any μ , and any $\mathcal{E} = (p^\theta)_{\theta \in \Theta} \in \Delta^o(S)^K$,

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

where $\omega(p^\theta) := (p_{\omega(1)}^\theta, \dots, p_{\omega(n)}^\theta)$ and $U_n(\cdot) = (U_n^1(\cdot), \dots, U_n^n(\cdot))$.

The final Axiom in this section requires that the updating is non-dogmatic and continuous. Continuity is satisfied by many models of belief revision in the literature, but may not hold if there are fixed costs of contemplation and belief revision; see Ortoleva (2012) for an example of discontinuous updating. Non-dogmatic revision, loosely stated, allows an agent to have arbitrary updated beliefs after a signal s if they observe sufficiently persuasive evidence. To be more precise, it says that if there are only two possible signals, then for any initial belief that attaches positive probability to every parameter value and any given updated belief after signal $s = 1$, there exists a unique experiment (or profile of probabilities for the signal $s = 1$) such that the updated belief after signal $s = 1$ has this value. This does *not* require that all possible profiles of beliefs can be generated from a suitable statistical experiment. It only that the range of updating function for a given signal is the entire set $\Delta^o(\Theta)$. The additional requirement that this map is a bijection is required for complete solution of the functional equation we later solve. Without uniqueness only local solutions for U_n exist.⁸ The uniqueness

⁷In 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 realised after the signal is observed. To capture this model of updating it would be necessary for the function U_n to take values in $\Delta(\Delta(\Theta))^n$.

⁸See, for example, Berg (1993).

property would again be violated by models updating that have fixed costs of contemplation. In such models there may be 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.

Axiom 3 (Non-Dogmatic). The function $U_2 : \Delta(\Theta) \times \Delta^o(S)^K \rightarrow \Delta(\Theta)^2$ is continuous. For any $\mu, \mu' \in \Delta^o(\Theta)$ there exists a unique $\tilde{\mathcal{E}} \in \Delta^o(S)^K$ such that $U_2^1(\mu, \tilde{\mathcal{E}}) = \mu'$.

2.1. Binary Experiments

We end this section by introducing the function u that describes the updating in the case where there are only two signals—a binary experiment. This will play a key role in what follows.

For a binary experiment with only two signals $p^\theta = (p_1^\theta, p_2^\theta)$ and $\mathcal{E} = (p_1^\theta, p_2^\theta)_{\theta \in \Theta}$ we can write

$$U_2(\mu, \mathcal{E}) \equiv (U_2^1(\mu, p_1, p_2), U_2^2(\mu, p_1, p_2)),$$

where $p_1 + p_2 = \mathbf{1}$ and $p_s = (p_s^1, \dots, p_s^K)$ are the probabilities of the signal $s \in \{1, 2\}$ for each of the parameters. Because of the axiom of symmetry the two functions on the right above are identical when the arguments are transposed, that is, $U_2^1(\mu, p_1, p_2) \equiv U_2^2(\mu, p_2, p_1)$. So, we define $u : \Delta(\Theta) \times (0, 1)^K \rightarrow \Delta(\Theta)$ as

$$(2) \quad u(\mu, p_1) := U_2^1(\mu, p_1, \mathbf{1} - p_1).$$

That is, $u(\mu, p_1)$ describes the updated beliefs after a binary experiment where the signal $s = 1$ occurred and p_1 were the parameter-dependent probabilities of the signal $s = 1$ and $\mathbf{1} - p_1$ were the probabilities for $s = 2$. This allows us to write the full profile of updated beliefs for binary experiments only in terms of one function u : and

$$U_2(\mu, \mathcal{E}) \equiv (u(\mu, p_1), u(\mu, p_2))$$

We will use divisibility to extend this decomposition so that it holds for experiments with arbitrary numbers of signals (4). Axiom 1 also implies that the function $u(\cdot)$ satisfies

$$(3) \quad u(\mu, p\mathbf{1}) = \mu, \quad \forall \mu \in \Delta(\Theta), p \in (0, 1).$$

3. AXIOMS ON ITERATED UPDATING: DIVISIBILITY AND ORDER INDEPENDENCE

We now consider two types of condition one might place on the updating of beliefs. The first condition on the updating we consider, Axiom 4, is termed divisibility. It is designed to capture the fact that the update is independent of how information is processed whether it is processed in an iterative way or as one-off process. The second condition we consider, Axiom 6 given in the Appendix, is that the updating of beliefs is independent of the order that information arrives. That is, reversing the order that two pieces of information arrive has no effect on the ultimate profile of beliefs. Although

these axioms appear different, we show in the Appendix that they are actually equivalent and so we will only use the divisibility axiom.

Consider two possible ways of learning the signal s . The first is a one-step process where a signal s is generated according to the experiment $\mathcal{E} = ((p^\theta)_{\theta \in \Theta})$ and then revealed to the individual. The second is a two-step process where, first, with probabilities $(p_1^\theta, 1 - p_1^\theta)_{\theta \in \Theta}$ the signal $s = 1$ or the signal $s \neq 1$ is revealed to the agent in a simpler 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 probabilities $(p_{-1}^\theta (1 - p_1^\theta)^{-1})_{\theta \in \Theta}$ (where p_{-s}^θ is the vector p^θ with the s^{th} element omitted). In what follows we will use $\mathcal{E}_{-1} := (p_{-1}^\theta (1 - p_1^\theta)^{-1})_{\theta \in \Theta}$ to denote the experiment that occurs conditional on $s \neq 1$.⁹

The Axiom 4 says that these two different processes of learning s have no effect on the individual's ultimate profile of beliefs. This assertion has two distinct elements. 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 and the other signal occurs with probabilities $1 - p_1^\theta$. Formally this requires that $U_n^1(\mu, \mathcal{E}) \equiv u(\mu, p_1)$.

The second element of the Axiom 4 says that the updated beliefs an individual has when they see the signal $s' > 1$ in a one-off experiment, that is $U_n^{s'}(\mu, \mathcal{E})$, are the same as the beliefs they would have in the two step-process. That is, if they were first told the signal was not $s = 1$ and updated their beliefs to $u(\mu, \mathbf{1} - p_1)$ and then were told the signal was $s' \in \{2, \dots, n\}$ from the experiment \mathcal{E}_{-1} and engaged in the further update to $U_{n-1}^{s'-1}(u(\mu, \mathbf{1} - p_1), \mathcal{E}_{-1})$.¹⁰ The form of this function requires a little explanation. The updated beliefs after the signal s' is the $s' - 1^{\text{th}}$ component of the vector U_{n-1} when the first signal is not present. Hence the change in the superscript on U_{n-1} .

Axiom 4 (Divisibility). For all $\mathcal{E} = ((p^\theta)_{\theta \in \Theta}) \in \Delta^o(S)^K$, $\mu \in \Delta(\Theta)$, and $n \geq 3$

$$U_n(\mu, \mathcal{E}) = [u(\mu, p_1), U_{n-1}(u(\mu, \mathbf{1} - p_1), \mathcal{E}_{-1})].$$

If this axiom is combined with symmetry and iteratively applied, it implies that any multi-step procedure that is based upon the experiment \mathcal{E} will result in the same updated profile of beliefs.

In the appendix we consider a second property for belief updating that might be desirable—that the order that information arrives has no effect on the ultimate profile of beliefs. That is, if signals arrive in an order and this order is re-arranged, then the ultimate profile of beliefs is unaffected by this alteration. We describe an axiom on the updating that generates this property and then show that this Axiom is equivalent to the divisibility axiom.

⁹Recall that $p^\theta \in \Delta^o(\Theta)$ and so $1 > p_1^\theta$.

¹⁰We use $\mathbf{1}$ to denote the vector $(1, 1, \dots, 1)$ of appropriate length.

4. CHARACTERISATION OF DIVISIBLE UPDATING

In this section a proposition is proved that gives a full characterisation of any updating procedure U_n that satisfies the Axioms 1, 2, 3, and 4. We will show that any such updating is characterised by a homeomorphism F that maps beliefs to shadow prior. Then, this shadow prior is updated in a fully Bayesian manner to create a shadow posterior. Finally the shadow posterior is mapped back by F^{-1} to form the individual's updated beliefs. This is illustrated in the figure below.




FIGURE 2. Updating Functions U that Satisfy Axioms 1, 2, 3, and 4.

The first step in the argument is to show that $u(\mu, p_s)$, the function that determines the updated beliefs after signal s , must be homogeneous degree zero in p_s if it satisfies Axioms 1, 2, and 4. The intuition for this result is quite simple. Suppose that there are three possible signals $S = \{1, 2, 3\}$ and that one signal, say $s = 2$, is equally likely under all states. We now apply the Axiom 4 to the two cases: one where s is determined in a one-off experiment and the second where there is first an experiment with binary outcomes $s = 2$ and $s \neq 2$ and then $s \in \{1, 3\}$ is selected with the appropriate conditional probabilities. As observing $s = 2$ is uninformative, Axiom 1 implies the first stage of the two-step process leads to no updating of the priors. However the second stage experiment, when a signal $s \neq 2$ is selected, has increased the relative probabilities p_s of the signals $s \in \{1, 3\}$. Thus the equality in Axiom 4 implies that the beliefs after the one-step experiment (with low relative probabilities p_s of the signals $s \in \{1, 3\}$) are equal to the beliefs after the two-step experiment (with higher relative probabilities of the signals). Thus scaling up the probabilities had no effect on the updating of beliefs. This is the definition of homogeneity degree zero.

Lemma 1. *Suppose updating satisfies Axioms 1, 2, and 4, then the function $u(\mu, p_s)$ defined in (2) is homogeneous degree zero in p_s . And*

$$(4) \quad U_n(\mu, \mathcal{E}) \equiv (u(\mu, p_1), \dots, u(\mu, p_n)).$$

Proof. Suppose that $n = 3$. Consider two ways the agent can process the signal $s = 1$: (a) She could be told the outcome of an $s = 1$ or $s > 1$ experiment, which would result in the updated beliefs (5). (b) The agent could first be told $s \neq 2$ in an $s = 2, s \neq 2$ experiment and then the outcome of a final experiment where $s \in \{1, 3\}$, which would result in the two step second updating (6).

$$(5) \quad u(\mu, p_1)$$

$$(6) \quad u\left(u(\mu, \mathbf{1} - p_2), \left(\frac{p_1^\theta}{1 - p_2^\theta}\right)_{\theta \in \Theta}\right)$$

The combination of Axioms 2 and 4 applied to (5) and (6) implies

$$(7) \quad u(\mu, p_1) \equiv u\left(u(\mu, \mathbf{1} - p_2), \left(\frac{p_1^\theta}{1-p_2^\theta}\right)_{\theta \in \Theta}\right) \quad \forall p_2^\theta \in [0, 1 - p_1^\theta].$$

If $1 - p_2^\theta = 1/\lambda$ for all θ , then (3) implies $u(\mu, \mathbf{1} - p_2) \equiv \mu$. So the right of (7) becomes $u(\mu, (\lambda p_1^\theta)_{\theta \in \Theta})$ and we get the condition

$$u(\mu, p_1) \equiv u(\mu, \lambda p_1), \quad \forall \lambda \in [1, \min_{\theta} (p_1^\theta)^{-1}].$$

Hence the function u is homogeneous degree zero in p_1 if Axiom 4 holds. The equation (4) follows by combining Axiom 4 with symmetry. \square

We now move on to establishing the main result which gives the form of the functions $u(\mu, p^s)$ defined in (2). The intuition for this characterisation of the function $u(\mu, p^s)$ comes from the fact that there many different intermediate experiments that will generate the same final updated beliefs. This is summarised in the functional equation, given below, that was derived in Lemma 1

$$u(\mu, p^s) \equiv u\left(u(\mu, \mathbf{1} - p_{s'}), \left(\frac{p_s^\theta}{1-p_{s'}^\theta}\right)_{\theta \in \Theta}\right) \quad \forall p_{s'}.$$

This says that the update for the signal s , $u(\mu, p^s)$, is also equal to the update for all equivalent intermediate experiments when the signal s' did not occur, then the individual had to do two updates. (They would first revise their beliefs to $u(\mu, \mathbf{1} - p_{s'})$ when s' didn't occur. Then, they would do a second update upon observing s .)

One could turn this equation into a family of PDE's and thereby determine u . However, there are techniques to solve this particular functional equation without the assumption of differentiability—continuity is enough. The first step is to turn it into a linear functional equation by taking logarithms. Giving the equation $u(\mu, p^s) \equiv u(u(\mu, \mathbf{1} - p_{s'}), p_s - (\mathbf{1} - p_{s'}))$. One simple solution to this equation is to add the arguments together: $u(\mu, p) = \mu + p$. Suitably adapted to the fact that μ , p^s , and $u(\cdot)$ are vectors of probabilities and the homogeneity result above, this simple solution gives Bayes' updating. For a more general solution one can first notice that the functional equation tells us about the contours of the function $u(\cdot)$. This is because as $p_{s'}$ varies the arguments on right, that is $u(\mu, \mathbf{1} - p_{s'})$ and $p_s - (\mathbf{1} - p_{s'})$ vary in a way that does not change the value of u . It is, then, relatively simple to see that each such contour of the function $u(\cdot)$ is a translation of the other. Thus once the form of one contour has been determined all other contours are just translations of it. Hence choosing an arbitrary function to determine the shape of one contour and another to determine the value taken by each contour is sufficient to determine the entire function. This is how the family of solutions to the functional equation given in the Proposition are generated.

The role of Axiom 3 here is to ensure that we are looking for a continuous solution to the functional equation, and that it holds globally not just locally. If continuity fails the set of solutions to even the simplest functional equation becomes enormous.¹¹ The role of uniqueness is to ensure a global solution exists, F is a homeomorphism, otherwise there may be many local solutions to the functional equation and only local solutions

¹¹See Aczél and Dhombres (1989) Chapter 1 for an example of this.

exist. These local solutions also have the same form as the global solution, see Berg (1993).

Proposition 1. If the updating U_n satisfies the Axioms 1, 2, 3, and 4, then there exists a homeomorphism $F : \Delta(\Theta) \rightarrow \Delta(\Theta)$ such that

$$(8) \quad u(\mu, p_s) = F^{-1} \left(\frac{F_1(\mu)p_s^1}{\sum_{\theta \in \Theta} F_\theta(\mu)p_s^\theta}, \dots, \frac{F_K(\mu)p_s^K}{\sum_{\theta \in \Theta} F_\theta(\mu)p_s^\theta} \right);$$

where $F(\mu) \equiv (F_1(\mu), F_2(\mu), \dots, F_K(\mu))$, and

$$U_n(\mu, \mathcal{E}) = (u(\mu, p_1), \dots, u(\mu, p_n)).$$

Proof. We begin by reducing the dimension of the variables in the function $u(\mu, p_s)$. Define $w : \Delta^\circ(\Theta) \rightarrow \mathbb{R}_{++}^{K-1}$ as follows

$$w(\mu_1, \dots, \mu_K) := \left(\frac{\mu_1}{\mu_K}, \dots, \frac{\mu_{K-1}}{\mu_K} \right);$$

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

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

We will now re-define the variables of the function u . First we define the variable $\phi := w(\mu)$, which requires us to restrict $\mu \in \Delta^\circ(\Theta)$ and use continuity to define $u(\mu, p_s)$ when μ is on the boundary of this set. We will also define $\pi := w(p_s)$. The fact $u(\mu, p_s)$ is homogeneous degree zero in $p_s \in (0, 1)^K$, by Lemma 1, implies that there is no loss by this transformation. We will also apply the transformation w to the function and define $v(\phi, \pi) \equiv w(u(\mu, p_s))$. Hence we have transformed a function $u : \Delta^\circ(\Theta) \times \Delta^\circ(\Theta) \rightarrow \Delta^\circ(\Theta)$ to a function $v : \mathbb{R}_{++}^{K-1} \times \mathbb{R}_{++}^{K-1} \rightarrow \mathbb{R}_{++}^{K-1}$.

If we re-write (7) with this new notation we get

$$v(\phi, \pi) \equiv v(v(\phi, \rho), \pi/\rho), \quad v : \mathbb{R}_{++}^{K-1} \times \mathbb{R}_{++}^{K-1} \rightarrow \mathbb{R}_{++}^{K-1};$$

where $\rho := w(\mathbf{1} - p_2)$ and $\pi/\rho := w((\frac{p_1^\theta}{1-p_2^\theta})_{\theta \in \Theta})$.

Now we take logarithms to do a final transformation of this function. Let us define $\tilde{\phi} := \ln \phi$, $\tilde{\pi} := \ln \pi$, $\tilde{\rho} = \ln \rho$, and $\tilde{v}(\tilde{\phi}, \tilde{\pi}) \equiv \ln v(\phi, \pi)$. Then, a final re-writing of (7) with this new notation gives

$$\tilde{v}(\tilde{\phi}, \tilde{\pi}) \equiv \tilde{v}(\tilde{v}(\tilde{\phi}, \tilde{\rho}), \tilde{\pi} - \tilde{\rho}), \quad \tilde{v} : \mathbb{R}^{K-1} \times \mathbb{R}^{K-1} \rightarrow \mathbb{R}^{K-1}.$$

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

$$\tilde{v}(\tilde{\phi}, y + z) \equiv \tilde{v}(\tilde{v}(\tilde{\phi}, z), y), \quad \forall \phi, y, z \in \mathbb{R}^{K-1}.$$

This is called the translation equation and was originally solved in its multivariate form by Aczél and Hosszú (1956). Given the uniqueness property in the regularity assumption, Axiom 3, the results described in Moszner (1995) page 21 apply. Thus any continuous solution to this equation has the property that there exists a continuous bijection $g : \mathbb{R}^{K-1} \rightarrow \mathbb{R}^{K-1}$ such that

$$(9) \quad \tilde{v}(\tilde{\phi}, \tilde{\pi}) = g^{-1}[g(\tilde{\phi}) + \tilde{\pi}].$$

Now we will reverse the transformations of this problem. First we will remove the logarithms in (9) to get

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

Then we will invert the g^{-1} function

$$g \circ \ln v(\phi, \pi) = g(\ln \phi) + \ln \pi.$$

Now we will introduce the function $h(x) = g(\ln x) = (h_1(x), \dots, h_{K-1}(x))$ to simplify this expression.

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

Finally, we define $f(x) \equiv e^{h(x)}$ and this becomes

$$f(v(\phi, \pi)) = (f_1(\phi) \pi_1, \dots, f_{K-1}(\phi) \pi_{K-1}).$$

Now substitute the $w(\cdot)$ transformation to get

$$(10) \quad f \circ w \circ u(\mu, p_s) = \left(f_1(w(\mu)) \frac{p_s^1}{p_s^K}, \dots, f_{K-1}(w(\mu)) \frac{p_s^{K-1}}{p_s^K} \right).$$

We will now define the function $F : \Delta^\circ(\Theta) \rightarrow \Delta^\circ(\Theta)$ so that the following diagram commutes, that is, $f \circ w \equiv w \circ F$. This is possible as w is invertible.

$$\begin{array}{ccc} \Delta^\circ(\Theta) & \xrightarrow{F} & \Delta^\circ(\Theta) \\ \downarrow w & & \downarrow w \\ \mathbb{R}_{++}^{K-1} & \xrightarrow{f} & \mathbb{R}_{++}^{K-1} \end{array}$$

F is a bijection and continuous hence it is a homeomorphism. We will extend this definition of F where necessary to the boundary of $\Delta(\Theta)$ using continuity. Using this we can re-write (10) as

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

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

$$F \circ u(\mu, p_s) \equiv \left(\frac{F_1(\mu) p_s^1}{\sum_{\theta \in \Theta} F_\theta(\mu) p_s^\theta}, \dots, \frac{F_K(\mu) p_s^K}{\sum_{\theta \in \Theta} F_\theta(\mu) p_s^\theta} \right).$$

Applying F^{-1} to both sides of this gives equation (8) in the proposition. The other displayed equation in the proposition follows from a substitution of (8) into (4) \square

As an example of this updating rule in action, consider again the model of the arrival of busses in Section 1.1. Recall that μ was the probability the individual attached to the good state and that a Bayesian updater would have a posterior $\frac{\alpha^t \mu}{\alpha^t \mu + \beta^t (1-\mu)}$ if there had been t periods without a bus arriving.

An individual who updates divisibly is described by a homeomorphism $F : (\mu, 1 - \mu) \mapsto (F_1(\mu), F_2(1 - \mu))$ and its inverse $F^{-1} : (\mu, 1 - \mu) \mapsto (F_1^{-1}(\mu), F_2^{-1}(1 - \mu))$. If such an individual arrived at the bus stop and learnt that there had been t periods without a bus arriving, then they would update their belief about the good state to

$$\mu_t = F_1^{-1} \left(\frac{F_1(\mu)\alpha^t}{F_1(\mu)\alpha^t + F_2(1 - \mu)\beta^t} \right).$$

One more period without a bus would then lead the individual to do the further update

$$\mu_{t+1} = F_1^{-1} \left(\frac{F_1(\mu_t)\alpha}{F_1(\mu_t)\alpha + F_2(1 - \mu_t)\beta} \right).$$

However, the updating rule for μ_t given here implies that

$$F_1(\mu_t) = \frac{F_1(\mu)\alpha^t}{F_1(\mu)\alpha^t + F_2(1 - \mu)\beta^t}, \quad \text{and} \quad F_2(1 - \mu_t) = 1 - F_1(\mu_t).$$

If these are substituted into the expression for μ_{t+1} we get

$$\mu_{t+1} = F_1^{-1} \left(\frac{F_1(\mu)\alpha^{t+1}}{F_1(\mu)\alpha^{t+1} + F_2(1 - \mu)\beta^{t+1}} \right).$$

Thus the beliefs μ_{t+1} of our individual are exactly those of a new arrival at the bus stop.

We now give two general classes of homeomorphisms F that generate particular classes of divisible updating rules.

4.1. Geometric Weighting

The first example of a homeomorphism $F : \Delta(\Theta) \rightarrow \Delta(\Theta)$ we consider results in an updating function that has already been used to model overreaction or under-reaction to new information.

$$F(\mu) = \left(\frac{\mu_1^{\alpha_1}}{\sum_{\theta} \mu_{\theta}^{\alpha_{\theta}}}, \dots, \frac{\mu_K^{\alpha_K}}{\sum_{\theta} \mu_{\theta}^{\alpha_{\theta}}} \right);$$

An explicit form for F^{-1} only exists when $\alpha = \alpha_{\theta}$ (for all θ).

$$F^{-1}(x) = \left(\frac{\mu_1^{1/\alpha}}{\sum_{\theta} \mu_{\theta}^{1/\alpha}}, \dots, \frac{\mu_K^{1/\alpha}}{\sum_{\theta} \mu_{\theta}^{1/\alpha}} \right), \quad \text{for } \alpha = \alpha_{\theta}, \quad \forall \theta.$$

We can now simply substitute this functional form into (8) to generate a belief updating rule that satisfies our Axioms. In general this gives the updating rule

$$u \left(\mu, (p_s^{\theta})_{\theta \in \Theta} \right) \equiv F^{-1} \left(\frac{\mu_1^{\alpha_1} p_s^1}{\sum_{\theta \in \Theta} \mu_{\theta}^{\alpha_{\theta}} p_s^{\theta}}, \dots, \frac{\mu_K^{\alpha_K} p_s^K}{\sum_{\theta \in \Theta} \mu_{\theta}^{\alpha_{\theta}} p_s^{\theta}} \right).$$

When $\alpha = \alpha_{\theta}$ for all θ

$$u \left(\mu, (p_s^{\theta})_{\theta \in \Theta} \right) \equiv \left(\frac{\mu_1 (p_s^1)^{1/\alpha}}{\sum_{\theta \in \Theta} \mu_{\theta} (p_s^{\theta})^{1/\alpha}}, \dots, \frac{\mu_K (p_s^K)^{1/\alpha}}{\sum_{\theta \in \Theta} \mu_{\theta} (p_s^{\theta})^{1/\alpha}} \right).$$

From this expression it is clear that this updating rule exaggerates or alters the probabilities that enter into the normal Bayesian formula. This can be interpreted as over-reaction or under-reaction to new information. But, in Section 5.2, we will also show that this generates a bias in the updating.

When $\alpha \neq \alpha_\theta$, there is no explicit solution for the full updated distribution. However, it is possible to proceed in a slightly different way to understand how the relative probabilities are updated. First, write

$$F \circ u \left(\mu, (p_s^\theta)_{\theta \in \Theta} \right) \equiv \left(\frac{\mu_1^{\alpha_1} p_s^1}{\sum_{\theta \in \Theta} \mu_\theta^{\alpha_\theta} p_s^\theta}, \dots, \frac{\mu_K^{\alpha_K} p_s^K}{\sum_{\theta \in \Theta} \mu_\theta^{\alpha_\theta} p_s^\theta} \right)$$

and then define $(\mu'_1, \dots, \mu'_K) = u \left(\mu, (p_s^\theta)_{\theta \in \Theta} \right)$ to get

$$\left(\frac{(\mu'_1)^{\alpha_1}}{\sum_{\theta} (\mu'_\theta)^{\alpha_\theta}}, \dots, \frac{(\mu'_K)^{\alpha_K}}{\sum_{\theta} (\mu'_\theta)^{\alpha_\theta}} \right) \equiv \left(\frac{\mu_1^{\alpha_1} p_s^1}{\sum_{\theta \in \Theta} \mu_\theta^{\alpha_\theta} p_s^\theta}, \dots, \frac{\mu_K^{\alpha_K} p_s^K}{\sum_{\theta \in \Theta} \mu_\theta^{\alpha_\theta} p_s^\theta} \right).$$

Then dividing the θ'' entry in this vector by the θ' entry we get

$$\frac{u_{\theta''}(\mu, (p_s^\theta)_{\theta \in \Theta})^{\alpha_{\theta''}}}{u_{\theta'}(\mu, (p_s^\theta)_{\theta \in \Theta})^{\alpha_{\theta'}}} = \frac{\mu_{\theta''}^{\alpha_{\theta''}} p_s^{\theta''}}{\mu_{\theta'}^{\alpha_{\theta'}} p_s^{\theta'}}; \quad \frac{u_{\theta''}(\mu, (p_s^\theta)_{\theta \in \Theta})}{u_{\theta'}(\mu, (p_s^\theta)_{\theta \in \Theta})} = \underbrace{\frac{\mu_{\theta''}}{\mu_{\theta'}} \left(\frac{p_s^{\theta''}}{p_s^{\theta'}} \right)^{1/\alpha}}_{\text{if } \alpha = \alpha_{\theta''} = \alpha_{\theta'}}.$$

In the case where the parameters α_θ differ it is still possible to give an explicit expression for the relative size of the updated belief in θ' and θ''

This updating rule generalises the overconfidence/under-confidence protocol described in Angrisani, Guarino, Jehiel, and Kitagawa (2017) and Bohren and Hauser (2017), for example. If $1/\alpha_\theta > 1$ then the individual overreacts to new information on parameter value θ and is overconfident by placing too much weight on their observations. This is achieved by exaggerating the differences in the signals and is more clearly seen in the ratio of the updated beliefs given above. Conversely, if $1/\alpha_\theta < 1$ then the individual under-reacts to new information about their learning on parameter value θ —they place too much weight on their prior and do not adjust their beliefs as much as a Bayesian would. In these two papers, the under and over-reaction is uniform across all parameters. But, the functional form permits the agent to be over-react for some parameters and under-react for others. Thus there is the possibility of selective over and under-reaction where the agent more readily changes beliefs about some parameters but not about others.

4.2. Exponential Weighting and Other Homeomorphisms

For all $\theta \in \Theta$, let $f_\theta : [0, 1] \rightarrow \mathbb{R}_+$ be a strictly increasing and continuous function that satisfies $f_\theta(0) = 0$. Then let us define F , as below

$$F(\mu) = \left(\frac{f_1(\mu_1)}{\sum_{\theta} f_\theta(\mu_\theta)}, \dots, \frac{f_K(\mu_K)}{\sum_{\theta} f_\theta(\mu_\theta)} \right).$$

Without an explicit form for F^{-1} we can use the relation

$$F \circ u \left(\mu, (p_s^\theta)_{\theta \in \Theta} \right) \equiv \left(\frac{f_1(\mu_1)p_s^1}{\sum_{\theta \in \Theta} f_\theta(\mu_\theta)p_s^\theta}, \dots, \frac{f_K(\mu_K)p_s^K}{\sum_{\theta \in \Theta} f_\theta(\mu_\theta)p_s^\theta} \right).$$

Again if we write the updated beliefs as $u \left(\mu, (p_s^\theta)_{\theta \in \Theta} \right) = (\mu'_1, \dots, \mu'_K)$, then the ratio of any two entries gives

$$\frac{f_\theta(\mu'_\theta)}{f_{\theta'}(\mu'_{\theta'})} = \frac{f_\theta(\mu_\theta)p_s^\theta}{f_{\theta'}(\mu_{\theta'})p_s^{\theta'}}.$$

As an example of such a function we could choose $f_\theta(\mu) = e^{-\beta_\theta/\mu}$, which results in a transformation of beliefs that is similar to a multinomial logit.

$$F(\mu) = \left(\frac{e^{-\beta_1/\mu_1}}{\sum_\theta e^{-\beta_\theta/\mu_\theta}}, \dots, \frac{e^{-\beta_K/\mu_K}}{\sum_\theta e^{-\beta_\theta/\mu_\theta}} \right) \quad \text{for } \beta_\theta > 0 \ \forall \theta.$$

This function F does not nest Bayesian updating—there are no values of the parameters β for which this function is the identity. Suppose that $\beta_\theta = \beta$ for all θ . Then F does map points closer to $\bar{\mu}$ when β is small and move points away from $\bar{\mu}$ when β is large. As $\beta \rightarrow 0$, so $F(\mu)$ converges to $\bar{\mu}$ for all interior μ . And as $\beta \rightarrow \infty$, so $F(\mu)$ converges to a distribution that puts all weight on the largest elements of μ . Thus the extremes of this function are similar to those in the previous example. And, using the intuitions from our previous example, we would expect large values of β to be associated with under-reaction and small values of β to be associated with over-reaction to new information.

An explicit form for the inverse function F^{-1} is not given here. However, by taking the ratio of any two entries (say $\hat{\theta}$ and $\tilde{\theta}$) in these vectors we get

$$\frac{e^{-\beta_{\hat{\theta}}/\mu'_{\hat{\theta}}}}{e^{-\beta_{\tilde{\theta}}/\mu'_{\tilde{\theta}}}} = \frac{e^{-\beta_{\hat{\theta}}/\mu_{\hat{\theta}}p_s^{\hat{\theta}}}}{e^{-\beta_{\tilde{\theta}}/\mu_{\tilde{\theta}}p_s^{\tilde{\theta}}}}.$$

Alternatively

$$\frac{\beta_{\tilde{\theta}}}{\mu'_{\tilde{\theta}}} - \frac{\beta_{\hat{\theta}}}{\mu'_{\hat{\theta}}} = \frac{\beta_{\tilde{\theta}}}{\mu_{\tilde{\theta}}} - \frac{\beta_{\hat{\theta}}}{\mu_{\hat{\theta}}} + \ln \frac{p_s^{\tilde{\theta}}}{p_s^{\hat{\theta}}}.$$

Thus the updating of beliefs results in a linear shift in the inverse probabilities of each parameter value. If $\beta_{\hat{\theta}} = \beta_{\tilde{\theta}} = \beta$ it is easy to see that large values of β decrease the dependence of this shift on on the ratio of the probabilities. Thus there is the conjectured under-reaction in this case.

5. PROPERTIES OF DIVISIBLE UPDATING

In this section we describe some of the general properties of the updating rules that satisfy the Axioms 1, 2, 3, and 4. In particular we will show that these updating rules obey consistency and describe conditions on the function F that ensure updating is biased in particular ways.

We are going to need to impose the condition that the updating respects certainty. That is, if the individual has a prior that attaches probability one to the parameter θ and observes signals that only occur with positive probability under this prior, then they do not revise their beliefs. We let $e_\theta \in \Delta(\Theta)$ denote the prior that attaches

probability one to the parameter θ , that is, it is a vector with unity in the θ entry and zeros elsewhere.

$$(11) \quad U_n(e_\theta, \mathcal{E}) = (e_\theta, \dots, e_\theta), \quad \text{for all } \mathcal{E} \in \Delta^o(S)^K.$$

5.1. The Consistency of Divisible Updating

Here we consider the limiting properties of the updating rule (8) as increasing amounts of information on the parameters $\theta \in \Theta$ is collected. For a given value of the parameter we will describe the limiting updated belief as the experiment \mathcal{E} is repeatedly sampled. We will show that the updating converges almost surely and that, if the experiment is informative, then the updated belief converges with probability one to the distribution e_θ —certainty that the parameter is θ (Proposition 2).¹² This result on consistency for divisible updating contrasts with examples of non-Bayesian updating in the literature that do not satisfy consistency: Rabin and Schrag (1999) and Epstein, Noor, and Sandroni (2010).¹³

In general models of learning there are two properties that might be desirable: (1) the individual attaches probability one to their beliefs converging, (2) the beliefs do converge with probability one. In general there are learning protocols where neither, all, or one of these properties hold. The convergence result below is conditional on the parameter value, thus, convergence independent of the initial beliefs the individual attaches to each possible parameter. Thus both of the above properties are satisfied by divisible learning.

We will begin with a description of the model where the individual repeatedly samples from the same experiment for a given parameter value. Fix a parameter value θ and let the stochastic process $\{s^t\}_{t=0}^\infty \in S^\infty$ be independently sampled from the distribution $p^\theta \in \Delta^o(S)$. We will use \mathbb{P}^θ to denote the probability measure on S^∞ of this process. We will also inductively define a belief stochastic process $\{\mu^t\}_{t=0}^\infty \in \Delta(\Theta)^\infty$ so that μ^{t+1} is the updated value of μ^t when the signal s^t is observed. The formal definition is: $\mu^0 \in \Delta(\Theta)$ and

$$(12) \quad \mu^{t+1} := u(\mu^t, p_{s^t}), \quad t = 0, 1, \dots$$

Proposition 2 shows that for all θ the process $\{\mu^t\}_{t=0}^\infty$ converges \mathbb{P}^θ almost surely provided the updating satisfies Axioms 1, 2, 3, 4. Furthermore, if the updating satisfies condition (11) and the signals are informative, then μ^t converges to e_θ , \mathbb{P}^θ almost surely. That is the updating satisfies consistency.

Proposition 2. If the updating U_n satisfies the Axioms 1, 2, 3, 4, then for all θ there exists $\mu^\infty \in \Delta(\Theta)$ such that $\mu^t \rightarrow \mu^\infty$, \mathbb{P}^θ almost surely. If the updating also respects certainty (11), $p^\theta \neq p^{\theta'}$ for all $\theta' \neq \theta$, and $\mu^0 \in \Delta^o(\Theta)$, then, $\mu^\infty = e_\theta$ with \mathbb{P}^θ probability one.

¹²This property of Bayesian updating is usually termed consistency, see Diaconis and Freedman (1986) for example.

¹³In Lehrer and Teper (2015) the notion of consistency is used as an axiom to characterise Bayesian updating.

The proof of this result is given in the appendix. The proof is quite trivial; it just applies the usual proof of the consistency of Bayesian updating (see for example DeGroot (1970)) to the shadow-belief revision process. It then uses the property that belief revision respects certainty (11) to ensure that when the shadow beliefs approach certainty they are mapped back (by a continuous inverse function) to beliefs that also approach certainty.

5.2. Bias in Divisible Updating

In this section we describe classes of divisible updating rules that exhibit biases for, or against, the truth. That is, we will give sufficient conditions on the function $F(\cdot)$ defined in (8) to generate such biases. To make the expression of these conditions simple we will consider a model with only two parameters $\Theta = \{\theta, \theta'\}$ where the individual forms the updated beliefs $u_\theta(\mu, p_s)$ when they see the signal s . We will say the updating is “locally consistent” if μ_θ is expected to increase when θ is true. This is the inequality

$$\mu_\theta \leq E^\theta(u_\theta(\mu, p_s)).$$

Here E^θ is an expectation taken relative to the objective distribution of the signals s . Thus, this inequality says that an outside (Bayesian) observer expects the individual to have an increased belief in θ when it is true and they observe the outcome of the experiment \mathcal{E} . Local consistency holds for Bayesian updating (which is globally consistent) and other models of updating that are approximately Bayesian.¹⁴ We will also say updating is “locally inconsistent” if μ_θ [respectively, $\mu_{\theta'}$] is expected to decrease [respectively, increase] when θ is true. These are the inequalities

$$\mu_\theta \geq E^\theta(u_\theta(\mu, p_s)), \quad \mu_{\theta'} \leq E^\theta(u_{\theta'}(\mu, p_s)).$$

This kind of bias does not arise here, because individuals are ignoring or misinterpreting their signals as in Rabin and Schrag (1999), for example. It arises because when θ is true the individual is slow to move their belief in θ upwards in response to positive evidence but quick to move beliefs down when evidence in favour of θ' is observed. These two effects on average give a downward movement of beliefs. Thus, the bias could be interpreted more as a reluctance to move to extreme beliefs rather than a universally positive bias in favour of parameter θ' . This negative movement in beliefs also appears to be at odds with the result in Proposition 2, that beliefs ultimately converge to the truth for all divisible updating processes. This conflict is resolved as we will show that an individual cannot be globally inconsistent, it is only possible for the individual to be locally inconsistent on a subset of the states.

In this two-parameter case the homeomorphism $F : (\mu_\theta, \mu_{\theta'}) \mapsto (\mu'_\theta, \mu'_{\theta'})$ can be described by its effect on its first element. We will use the strictly increasing function $f : [0, 1] \rightarrow [0, 1]$ to describe how F operates on its first element

$$F(\mu) \equiv (f(\mu_\theta), 1 - f(\mu_\theta)).$$

Before stating the result we will give a definition of *locally convex* and *locally concave* function. This allows us to think about the local biases in the updating rather than only

¹⁴This is a feature of some of the arguments in Bohren and Hauser (2017) for example.

global properties. We introduce the set $R(\mu_\theta)$ to describe all possible updated beliefs that the experiment \mathcal{E} could generate given the initial beliefs μ_θ and the updating function $u(\cdot)$. More precisely we let $R(\mu_\theta) \subset [0, 1]$ be the smallest interval that contains all the possible values for the updated beliefs when the initial belief is μ_θ .

$$R(\mu_\theta) := \left\{ \mu_\theta : \min_s f^{-1}\left(\frac{f(\mu_\theta)p_s^\theta}{f(\mu_\theta)p_s^\theta + (1-f(\mu_\theta))p_s^{\theta'}}\right) \leq \mu \leq \max_s f^{-1}\left(\frac{f(\mu_\theta)p_s^\theta}{f(\mu_\theta)p_s^\theta + (1-f(\mu_\theta))p_s^{\theta'}}\right) \right\}.$$

The definition of local concavity/convexity of a function g simply requires g to be concave/convex on the range of updated beliefs that could occur when the original belief is μ_θ , that is, the interval $R(\mu_\theta)$. If the experiment is particularly informative or if f increases rapidly this requires g to be concave/convex on a large interval but it also may be a relatively small interval for uninformative experiments or unresponsive functions f .

Definition 1. *The function $g : [0, 1] \rightarrow \mathbb{R}$ is locally convex at μ_θ , if it is a convex function on the set $R(\mu_\theta)$. The function $g : [0, 1] \rightarrow \mathbb{R}$ is locally concave at μ_θ , if it is a concave function on the set $R(\mu_\theta)$.*

Equipped with this definition we can provide sufficient conditions for the biases described above. The proof of this proposition is given in the appendix.

Proposition 3. *If the function f is increasing:*

- (i) *And $1/f(\cdot)$ is locally convex at μ_θ , then $\mu_\theta \leq E^\theta(u_\theta(\mu, p_s))$ and updating is locally consistent.*
- (ii) *And $1/f(\cdot)$ is locally concave at μ_θ , then $\mu_\theta \geq E^\theta(u_\theta(\mu, p_s))$ and updating is locally inconsistent.*

If $f(0) = 0$ then $1/f$ is not locally concave for all μ_θ .

For Bayesian updating $f(\mu)$ is the identity and in this case $1/f(\mu) = \mu^{-1}$ is convex. In this case Proposition 3 is the usual result that a Bayesian has beliefs that are a submartingale and on average are revised upwards.

Now we will revisit the example in Section 4.1 where F has the geometric weighting form (and $\alpha_\theta = \alpha_{\theta'} = \alpha$). When there are only two parameters this can be simplified to $f(\mu) = \frac{\mu^\alpha}{\mu^\alpha + (1-\mu)^\alpha}$ (this function is illustrated for $\alpha < 1$ in Figure 5.2 below). First observe that $1/f(\mu) = 1 + \left(\frac{1-\mu}{\mu}\right)^\alpha$ is convex when $\alpha \geq 1$. So, by case (i) of Proposition 3, updating is always locally consistent when $\alpha \geq 1$. However, when $\alpha < 1$ there exists an interval of values $\mu \in (\frac{1}{2}(1+\alpha), 1]$ where $1/f(\mu)$ is concave. In this region the updating is locally inconsistent (case (ii) of the Proposition). Thus as the individual gets closer to certain that θ is the true parameter, the expected movement of the beliefs is away from certainty. One could interpret this as a sort of mean reversion in the updating. The convergence to certainty is slower as a result. It is clear from the figure below that (when $\alpha < 1$) f tends to move beliefs to more central positions. As the experimenter becomes more convinced that the parameter is θ , $\mu \geq \frac{1}{2}(1+\alpha)$, the updating on average moves the updated beliefs away from these extreme values—revising beliefs downwards in response to bad news on θ much faster than they are increased in response to good news. Although, the bad news is less likely than the good news, the curvature of f

is sufficient to ensure that the expectation of these two effects is less than the current belief in θ .




FIGURE 3. $f(\mu) = \frac{\mu^\alpha}{\mu^\alpha + (1-\mu)^\alpha}$, for $\alpha < 1$.

5.3. Sufficient Conditions for Under and Overreaction to Information

In this section we provide a sufficient condition for the updating rule from Proposition 1 to overreact or under-react to new information.

We will argue that one cause of overreaction is that the homeomorphism F contracts space or moves points closer together. Once the shadow prior has been updated to a shadow posterior, the distance between the shadow prior and shadow posterior is a measure of how much Bayesian learning has occurred. If the inverse function F^{-1} expands space (or F contracts space), then when the shadow posterior 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 expands space and moves points further apart. Then the learning that occurred in the shadow Bayesian world gets reduced when it is mapped back to the belief space by F^{-1} , because the inverse will move points closer together. As a result the Bayesian learning in the shadow space is understated and there is under-reaction to new information.

This is not the complete explanation of the causes of under-reaction, because the amount of Bayesian learning that occurs is also influenced by the prior that is being updated. For example if the prior is close to certainty, then it will not be revised much in response to a given experiment. Thus under-reaction to an experiment can be generated if μ is mapped to an $F(\mu)$ that is close to certainty, because in this case the Bayesian update at $F(\mu)$ will be small. Thus there are two effects that influence under and over-reaction to information: how F expands or contracts space and the extent to which F maps μ to the extreme values.

There are homeomorphisms in which both these effects tend to work together. This is the case if F mapped all interior beliefs closer to the centre of $\Delta(\Theta)$. Then, the Bayesian update of the resultant shadow prior would be larger because the shadow

prior is more uncertain and less extreme. And the inverse map F^{-1} would expand space (as F contracts space) and move these Bayesian updates further apart.¹⁵

As in the previous section we are able to give sufficient conditions for over- and under-reaction for the case where there are two parameters, $\Theta = \{\theta, \theta'\}$. In this case we can treat the homeomorphism $F(\mu) \equiv (f(\mu_\theta), 1 - f(\mu_\theta))$ as being determined by a one-dimensional function $f : [0, 1] \rightarrow [0, 1]$.¹⁶

To define over-reaction or under-reaction it is necessary to do three things: first, select a measure of the beliefs, second choose a measure of variability of beliefs; and third specify a benchmark model of updating to make comparisons with. The results have a simpler form if we measure beliefs using the odds ratio $\frac{\mu_\theta}{\mu_{\theta'}} = \frac{\mu_\theta}{1 - \mu_\theta}$. The variance will be our measure of variability of beliefs and the Bayesian update as the benchmark for comparison. Thus our benchmark measure of variability is the variance of the odds ratio of the Bayesian update. If we use μ'_θ to denote the Bayesian update, then the odds ratio has the variance

$$\text{Var} \left[\frac{\mu'_\theta}{1 - \mu'_\theta} \right] = \text{Var} \left[\frac{\mu_\theta}{1 - \mu_\theta} \frac{p_s^\theta}{p_s^{\theta'}} \right] = \left(\frac{\mu_\theta}{1 - \mu_\theta} \right)^2 \text{Var} \left[\frac{p^\theta}{p^{\theta'}} \right].$$

The final variance is zero if and only if the signals in the experiment are uninformative $p^\theta = p^{\theta'}$. It is clear this benchmark is affect both by the information content of the experiment and by the initial prior.

Now we can define under-reaction to new information to be the situation where the variance of the odds ratio of the update is less than the variance of the Bayesian update and over-reaction to hold when it is greater.

Definition 2. *The updating exhibits under-reaction at μ if*

$$(13) \quad \text{Var} \left[\frac{u_\theta(\mu, p_s)}{1 - u_\theta(\mu, p_s)} \right] < \left(\frac{\mu_\theta}{1 - \mu_\theta} \right)^2 \text{Var} \left[\frac{p^\theta}{p^{\theta'}} \right],$$

and it exhibits overreaction at μ if

$$(14) \quad \text{Var} \left[\frac{u_\theta(\mu, p_s)}{1 - u_\theta(\mu, p_s)} \right] > \left(\frac{\mu_\theta}{1 - \mu_\theta} \right)^2 \text{Var} \left[\frac{p^\theta}{p^{\theta'}} \right].$$

We now are able to state a result on when divisible updating exhibits over and under-reaction in the two-parameter case.

Proposition 4. Suppose that the updating is described by the homeomorphism $F(\mu) \equiv (f(\mu_\theta), 1 - f(\mu_\theta))$, that f is differentiable, and that there exists a constant $\beta > 1$ such that

$$\beta \geq \frac{\mu_\theta}{f(\mu_\theta)} \geq \frac{1}{\beta}, \quad \beta \geq \frac{1 - \mu_\theta}{1 - f(\mu_\theta)} \geq \frac{1}{\beta};$$

for all $\mu_\theta \in [0, 1]$. Then: (i) if $\max_{\mu \in R(\mu_\theta)} f'(\mu) \leq \beta^{-4}$, then the updating exhibits overreaction at μ_θ . (ii) if $\min_{\mu \in R(\mu_\theta)} f'(\mu) \geq \beta^4$, then the updating exhibits under-reaction at μ_θ .

¹⁵An example of such a homeomorphism is given in Figure 5.2 above.

¹⁶The principal reason the general case of this result (for more than two parameters) is out of reach is that, apart from the linear case, there is no simple relationship between $\text{Var}[\mu]$ and $\text{Var}[F(\mu)]$ when μ is multidimensional.

The sufficient conditions for under and over-reaction depend on the two factors that we have already mentioned. The expansion/contraction of space is captured by the slope of the function $f(\mu)$. The difference between μ and $F(\mu)$ is bounded by the factor $\beta > 1$. The conditions below require the slope of f is greater than β^4 for under-reaction, or that the slope of f is less than β^{-4} for over-reaction. Thus it is indeed necessary for f to move points farther apart for an under-reaction to occur, but this effect is tempered by any movement that may have occurred in the mapping of beliefs to the shadow space.

6. CHARACTERISATION OF BAYESIAN UPDATING

In this section we show that one additional axiom—that the updating is unbiased, or a martingale—is sufficient for full Bayesian updating.

One might view of Proposition 1 as saying that the axioms we have provided so far are virtually enough for Bayesian updating. Thus it seems likely that a small additional requirement will give characterise Bayesian updating. We do not claim that the martingale property is minimal in this sense. There may well be weaker restrictions on the updating process that when added to Axioms 1, 2, 3, and 4 give Bayesianism. However, the unbiased nature of belief revision is such a fundamental property that it seems important to consider it directly. Furthermore the non-Bayesian updating of Epstein, Noor, and Sandroni (2010) is a martingale. Hence, it provides a useful example of a belief revision process that satisfies Axioms 1, 2, 3 and the martingale condition but is not Bayesian.¹⁷ Thus Axiom 4 is essential for these results.

Another route to take in this section would be to try to find an axiom that rules out features of updating that are non-Bayesian but are nevertheless consistent with Proposition 1. For example, to try to find an axiom that does not admit over-reaction or under-reaction to new information. This is not what the martingale axiom does. Indeed the model of Epstein, Noor, and Sandroni (2010) does permit over-reaction or under-reaction and is also a martingale. Thus it appears that it is the interaction of the martingale property and the divisibility property that jointly act to give Bayesian updating.

The Axiom below considers the profile of updated belief distributions: this is the distribution $U_n^s(\mu, (p)_{\theta \in \Theta}) \in \Delta(\Theta)$ for each signal s . Then it averages these distributions with the weights $\sum_{\theta \in \Theta} \mu_{\theta} p_s^{\theta}$, which are the ex-ante probabilities of the signals $s \in S$. This average is the decision maker’s predicted distribution of updated beliefs and the axiom requires that this predicted distribution equals the original beliefs. This property is sometimes termed “Bayes plausibility”, but can be required of non-Bayesian updating also.

¹⁷This revision process only satisfies Axiom 3 when there is over-reaction. In the case of underreaction extreme updated beliefs are not possible.

Axiom 5 (Unbiased). For any $\mu > 0$, $n > 1$, and $\mathcal{E} \in \Delta^o(S)^K$

$$\mu \equiv \sum_{s \in S} \left(\sum_{\theta \in \Theta} \mu_\theta p_s^\theta \right) U_n^s(\mu, (p)_{\theta \in \Theta}).$$

We will now prove the following result

Proposition 5. If the updating U_n satisfies the Axioms 1, 2, 3, 4, and 5, then it satisfies Bayes rule, that is,

$$(15) \quad u(\mu, p_s) \equiv \left(\frac{\mu p_s^1}{\sum_{\theta \in \Theta} \mu p_s^\theta}, \dots, \frac{\mu p_s^K}{\sum_{\theta \in \Theta} \mu p_s^\theta} \right).$$

and

$$U_n(\mu, \mathcal{E}) \equiv (u(\mu, p_1), \dots, u(\mu, p_K)).$$

Proof. It is sufficient to prove that the function $F(\cdot)$ in Proposition 1 is the identity. Without loss, we will consider the case where $n = 2$ and there is a binary experiment with signals $S = \{1, 2\}$. When the parameter is θ we will write the probabilities of the two signals as $(p^\theta, 1 - p^\theta)$.

In this case the unbiased condition, Axiom 5, is equivalent to:

$$(16) \quad \mu \equiv \left(\sum_{\theta} \mu_\theta p^\theta \right) F^{-1}(u_1^1, \dots, u_K^1) + \left(\sum_{\theta} \mu_\theta (1 - p^\theta) \right) F^{-1}(u_1^2, \dots, u_K^2),$$

$$\text{where } u_{\theta'}^1 = \frac{F_{\theta'}(\mu) p^{\theta'}}{\sum_{\theta} F_{\theta}(\mu) p^\theta}, \quad u_{\theta'}^2 = \frac{F_{\theta'}(\mu) (1 - p^{\theta'})}{\sum_{\theta} F_{\theta}(\mu) (1 - p^\theta)}, \quad \forall \theta' \in \Theta.$$

This holds for all values of $\mu \in \Delta^o(\Theta)$, all $p^\theta \in [0, 1]$, and all $q^\theta \in [0, 1]$, because $F(\cdot)$ is continuous and also defined on its boundary.

We will begin by showing that F is the identity at the extreme points of the set $\Delta(\Theta)$. Letting $p^{\theta'} \rightarrow 0$ in (16) for all $\theta' \neq \theta$ and $p^\theta \rightarrow 1$ gives the equality

$$\mu = \mu_\theta F^{-1}(e_\theta) + (1 - \mu_\theta) F^{-1}(y_\theta).$$

where: e_θ is a vector with 1 in the θ element and zeros elsewhere and y_θ is a vector with zero in the θ element and $F_{\theta'}(\mu)/(1 - F_\theta(\mu))$ for $\theta' \neq \theta$. The θ component of this vector equality is

$$\frac{\mu_\theta}{1 - \mu_\theta} (1 - F_\theta^{-1}(e_\theta)) = F_\theta^{-1}(y_\theta).$$

If the parenthesis is non-zero, the left of this expression can be arbitrarily large as μ_θ varies. But the right takes values only in $[0, 1]$. Thus, the only way this condition can hold for all $\mu_\theta \in (0, 1)$ is if: $1 = F_\theta^{-1}(e_\theta)$ and $0 = F_\theta^{-1}(y_\theta)$. As $F^{-1}(\mu) \in \Delta(\Theta)$, the first of these implies that $e_\theta = F^{-1}(e_\theta)$ and F is the identity on its extreme points. The second condition, $0 = F_\theta^{-1}(y_\theta)$, implies that F maps zero probabilities to zero probabilities.

We now derive a functional equation (18) which will allow us to establish the Proposition. Letting $p^{\theta'} \rightarrow 0$ for all $\theta' \neq \theta$ in (16) and imposing the continuity of F now implies

$$\mu \equiv \mu_\theta p^\theta F^{-1}(e_\theta) + (1 - \mu_\theta p^\theta) F^{-1} \left(\frac{F(\mu) - p^\theta F_\theta(\mu) e_\theta}{1 - p^\theta F_\theta(\mu)} \right).$$

As $F^{-1}(e_\theta) = e_\theta$ this rearranges to give

$$(17) \quad F\left(\frac{\mu - p^\theta \mu_\theta e_\theta}{1 - \mu_\theta p^\theta}\right) \equiv \frac{F(\mu) - p^\theta F_\theta(\mu) e_\theta}{1 - p^\theta F_\theta(\mu)}.$$

We will reduce the dimension of the function F by dropping one dimension μ_θ of its domain. Observe that $\mu_{-\theta} \in \mathcal{S} := \{x \in \mathbb{R}_+^{K-1} : \mathbf{1}^T x \leq 1\}$. Now define $G_\theta : \mathcal{S} \rightarrow \mathcal{S}$, so that $G_\theta(\mu_{-\theta})$ describes how F maps the vector $\mu_{-\theta}$ into \mathcal{S} , as follows

$$F(\mu) \equiv (G_\theta(\mu_{-\theta}), 1 - \mathbf{1}^T G_\theta(\mu_{-\theta})).$$

(As F is a homeomorphism, so too is G .) Dropping the θ th row in the vector equation (17) and re-writing what remains using the G_θ notation then gives

$$G_\theta\left(\frac{\mu_{-\theta}}{1 - \mu_\theta p^\theta}\right) \equiv \frac{G_\theta(\mu_{-\theta})}{1 - p^\theta F_\theta(\mu)}.$$

Finally, substituting $\lambda = (1 - \mu_\theta p^\theta)^{-1} \in [1, (1 - \mu_\theta)^{-1}]$ into the above then gives

$$(18) \quad G_\theta(\lambda \mu_{-\theta}) \equiv \frac{\lambda}{\lambda - (1 - \lambda) \frac{F_\theta(\mu)}{\mu_\theta}} G_\theta(\mu_{-\theta}), \quad \forall \lambda \in [1, (1 - \mu_\theta)^{-1}].$$

Thus the function G_θ satisfies a property that is similar to homogeneity (full homogeneity holds when $F_\theta(\mu) = \mu_\theta$).

For any $\mu \in \Delta^\circ(\Theta)$ consider the sequence $\{\mu^n\}_{n=1}^\infty$ where

$$\mu^n := \left(\frac{\mu_{-\theta}}{n}, \frac{n-1 + \mu_\theta}{n}\right).$$

Along this sequence we have that $\mu_\theta^n \rightarrow 1$ and $F_\theta(\mu^n) \rightarrow 1$ by the previously established result ($F(e_\theta) = e_\theta$) and the continuity of F . Writing the relation (18) for μ_n gives

$$G_\theta\left(\frac{\lambda'}{n} \mu_{-\theta}\right) \equiv \frac{\lambda'}{\lambda' - (1 - \lambda') \frac{F_\theta(\mu^n)}{\mu_\theta^n}} G_\theta\left(\frac{1}{n} \mu_{-\theta}\right), \quad \forall 1 \leq \lambda' \leq \frac{n}{1 - \mu_\theta}.$$

If we make the choices $\lambda' = n$ and $\lambda' = n\lambda$, then we can use the left of this expression to substitute for the terms $G_\theta(\lambda \mu_{-\theta})$ and $G_\theta(\mu_{-\theta})$ in (18). This then gives

$$\frac{n}{n - (1 - n) \frac{F_\theta(\mu^n)}{\mu_\theta^n}} G_\theta\left(\frac{1}{n} \mu_{-\theta}\right) = \frac{\lambda}{\lambda - (1 - \lambda) \frac{F_\theta(\mu)}{\mu_\theta}} \frac{n\lambda}{n\lambda - (1 - n\lambda) \frac{F_\theta(\mu^n)}{\mu_\theta^n}} G_\theta\left(\frac{1}{n} \mu_{-\theta}\right)$$

Hence for all n

$$1 = \frac{\lambda}{\lambda - (1 - \lambda) \frac{F_\theta(\mu)}{\mu_\theta}} \frac{\lambda - \lambda(\frac{1}{n} - 1) \frac{F_\theta(\mu^n)}{\mu_\theta^n}}{\lambda - (\frac{1}{n} - \lambda) \frac{F_\theta(\mu^n)}{\mu_\theta^n}}$$

As we know that $\mu_\theta^n \rightarrow 1$ and $F_\theta(\mu^n) \rightarrow 1$ this implies that

$$1 = \frac{\lambda}{\lambda - (1 - \lambda) \frac{F_\theta(\mu)}{\mu_\theta}} \quad \forall \lambda \in [1, (1 - \mu_\theta)^{-1}]$$

which can only be true if $F_\theta(\mu) = \mu_\theta$. This holds for all θ and all $\mu \in \Delta^\circ(\Theta)$, so $F(\mu) = \mu$ for all $\mu \in \Delta^\circ(\Theta)$. As F is continuous on $\Delta(\Theta)$ it follows that F is the identity. Substitution of $F(\mu) = \mu$ into (8) establishes the claim in this proposition. \square

7. EXTENSIONS AND APPLICATIONS

The assumption that the range of the updated beliefs is the whole probability simplex $\Delta(\Theta)$, made here, is inconsistent with several important models of non-Bayesian belief updating. For example, if the individual has limited memory or mental capacities, one might want to consider updating procedures that generate one of a finite set of probability measures. Thus the range of the updating function is finite. Models of learning where updated beliefs are discrete and finite, such as Hellman and Cover (1970), Dow (1991), and Wilson (2014), pose a considerable challenge for divisibility. It is not clear whether it is ever possible to satisfy divisibility of belief revision in such a setting.

There are other models of belief revision where the updates themselves are random. Thus the updating function maps an experiment and an initial belief to a profile of probability measures over updated beliefs. One example of this random updating is Rabin and Schrag (1999). It seems conceivable that a generalised notion of divisibility might apply in this setting.

REFERENCES

- ACZÉL, D., AND J. DHOMBRES (1989): *Functional Equations in Several Variables*. Cambridge University Press, Cambridge, UK, second edn.
- 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.
- BERG, L. (1993): “The Local Structure of the Solutions of the Multidimensional Translation Equation,” *Aequationes Mathematicae*, 46, 164–173.
- BOHREN, A., AND D. HAUSER (2017): “Bounded Rationality and Learning, A Framework and a Robustness Result,” *under review*, pp. 349–374.
- BRUNNERMEIER, M. K. (2009): “Deciphering the Liquidity and Credit Crunch 2007–2008,” *Journal of Economic Perspectives*, 23(1), 77–100.
- DAWID, A. P. (1984): “Present position and potential developments: Some personal views: Statistical theory: The prequential approach,” *Journal of the Royal Statistical Society. Series A (General)*, 147(2), 278–292.
- 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.
- DOW, J. (1991): “Search decisions with limited memory,” *The Review of Economic Studies*, 58(1), 1–14.

- EPSTEIN, L. G., J. NOOR, AND A. SANDRONI (2010): “Non-Bayesian Learning,” *The B.E. Journal of Theoretical Economics*, 10(1).
- EPSTEIN, L. G., AND M. SCHNEIDER (2003): “Recursive Multiple-Priors,” *Journal of Economic Theory*, 113(1), 1–31.
- EPSTEIN, L. G., AND S. E. ZIN (1989): “Substitution, Risk Aversion, and the Temporal Behavior of Consumption and Asset Returns: A Theoretical Framework,” *Econometrica*, 57(4), 937–969.
- GILBOA, I., AND D. SCHMEIDLER (1993): “Updating ambiguous beliefs,” *Journal of economic theory*, 59(1), 33–49.
- 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.
- HELLMAN, M. E., AND T. M. COVER (1970): “Learning with finite memory,” *The Annals of Mathematical Statistics*, pp. 765–782.
- KALLENBERG, O. (2002): *Foundations of Modern Probability*. Springer Verlag, New York, second edn.
- 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.
- MOSZNER, Z. (1995): “General Theory of the Translation Equation,” *Aequationes Mathematicae*, 50, 17–37.
- ORTOLEVA, P. (2012): “Modeling the change of paradigm: Non-Bayesian reactions to unexpected news,” *American Economic Review*, 102(6), 2410–36.
- RABIN, M., AND J. L. SCHRAG (1999): “First Impressions Matter: A Model of Confirmatory Bias,” *The Quarterly Journal of Economics*, 114(1), 37–82.
- SHANNON, C. E., AND W. WEAVER (1949): *The Mathematical Theory of Communication*. University of Illinois Press, Urbana and Chicago.
- TORGERSEN, E. (1991): *Comparison of Statistical Experiments*. Cambridge University Press, Cambridge, UK.
- TVERBERG, H. (1958): “A New Derivation of the Information Function,” *Mathematica Scandinavica*, 6, 297–298.
- WILSON, A. (2014): “Bounded memory and biases in information processing,” *Econometrica*, 82(6), 2257–2294.
- ZHAO, C. (2016): “Pseudo-Bayesian Updating,” mimeo.

APPENDIX

The Calculations for the Example in Section 1.1

Recall that $\mu = 1/2$ and:

$$\begin{aligned}\tilde{\mu}_2 &= (1 - \lambda)\frac{1}{2} + \lambda\frac{\alpha^2}{\beta^2 + \alpha^2}, \\ \mu_1 &= (1 - \lambda)\frac{1}{2} + \lambda\frac{\alpha}{\beta + \alpha}, \quad \mu_2 = (1 - \lambda)\mu_1 + \lambda\frac{\alpha\mu_1}{(1 - \mu_1)\beta + \mu_1\alpha}.\end{aligned}$$

We now take this expression for $\tilde{\mu}_2$ and make a sequence of substitutions. First for μ_1 and then for μ_2 . A final re-arranging then gives the displayed equation.

$$\begin{aligned}\tilde{\mu}_2 &= (1 - \lambda)\frac{1}{2} + \lambda\frac{\alpha^2}{\beta^2 + \alpha^2} \\ &= \mu_1 + \lambda\frac{\alpha^2}{\beta^2 + \alpha^2} - \lambda\frac{\alpha}{\beta + \alpha} \\ &= \mu_2 + \lambda\frac{\alpha^2}{\beta^2 + \alpha^2} - \lambda\frac{\alpha}{\beta + \alpha} + \lambda\mu_1 - \lambda\frac{\alpha\mu_1}{(1 - \mu_1)\beta + \mu_1\alpha} \\ \frac{\tilde{\mu}_2 - \mu_2}{\lambda} &= -\frac{\alpha\beta(\beta - \alpha)}{(\beta^2 + \alpha^2)(\beta + \alpha)} + \frac{(\beta - \alpha)\mu_1(1 - \mu_1)}{(1 - \mu_1)\beta + \mu_1\alpha} \\ \frac{\tilde{\mu}_2 - \mu_2}{\lambda(\beta - \alpha)} &= -\frac{\alpha\beta}{(\beta^2 + \alpha^2)(\beta + \alpha)} + \frac{1}{\frac{\beta}{\mu_1} + \frac{\alpha}{1 - \mu_1}}\end{aligned}$$

The final term above is an increasing function of μ_1 when $\mu_1 \in [0, 1/2]$. The RHS is positive when $\mu_1 = 1/2$ and negative when $\mu_1 = 0$. As λ increases from zero, so μ_1 decreases from $\frac{1}{2}$ to zero. Thus there exists a unique value $\bar{\lambda}$ such that $\tilde{\mu}_2 > \mu_2$ if and only if $\lambda < \bar{\lambda}$. But $\tilde{\mu}_2 = \mu_2$ when $\lambda = 1$ so we know $\tilde{\mu}_2 > \mu_2$ if and only if $\lambda < 1$.

Continuous Limit of the Example in Section 1.1

Writing the updating of μ_t in terms of the odds ratios we get

$$(19) \quad \frac{\mu_{t+1}}{1 - \mu_{t+1}} = \frac{\mu_t}{1 - \mu_t} \left(1 - \frac{\lambda(\beta - \alpha)}{\beta - \mu_t(\beta - \alpha)(1 - \lambda)} \right).$$

We have the two methods of updating beliefs in this example:

$$\begin{aligned}\frac{\mu_t}{1 - \mu_t} &= \prod_{s=0}^{t-1} \left(1 - \frac{\lambda(\beta - \alpha)}{\beta - \mu_s(\beta - \alpha)(1 - \lambda)} \right) \frac{\mu}{1 - \mu}, \\ \frac{\tilde{\mu}_t}{1 - \tilde{\mu}_t} &= \left(1 - \frac{\lambda(\beta^t - \alpha^t)}{\beta^t - \mu(\beta^t - \alpha^t)(1 - \lambda)} \right) \frac{\mu}{1 - \mu}.\end{aligned}$$

We will consider what happens as the period length shrinks: $dt \rightarrow 0$, the number of periods in any time interval $[0, \tau]$ tends to infinity $t = \tau/dt \rightarrow \infty$, and the arrival rate of busses also shrinks: $1 - \alpha = adt$, $1 - \beta = bdt$. This implies $\alpha^t = (1 - adt)^{\tau/dt} \rightarrow e^{-a\tau}$

and $\beta^t = (1 - bdt)^{\tau/dt} \rightarrow e^{-b\tau}$, hence

$$\frac{\tilde{\mu}_\tau}{1 - \tilde{\mu}_\tau} := \left(1 - \frac{\lambda(1 - e^{-(a-b)\tau})}{1 - \mu(1 - e^{-(a-b)\tau})(1 - \lambda)} \right) \frac{\mu}{1 - \mu}.$$

We have the recursion for the iterated updating:

$$\frac{\mu_{t+dt}}{1 - \mu_{t+dt}} = \frac{\mu_t}{1 - \mu_t} (1 - \lambda(a - b)dt).$$

Letting $dt \rightarrow 0$ and solving the resultant differential equation gives

$$\frac{\mu_\tau}{1 - \mu_\tau} = \frac{\mu}{1 - \mu} e^{-\lambda(a-b)\tau}.$$

These imply that:

$$\mu_\tau = \frac{e^{-\lambda(a-b)\tau} \mu}{1 - \mu + \mu e^{-\lambda(a-b)\tau}}, \quad \tilde{\mu}_\tau = (1 - \lambda)\mu + \lambda \frac{e^{-(a-b)\tau} \mu}{1 - \mu + \mu e^{-(a-b)\tau}},$$

Letting $\psi := a - b$, we have that $\tilde{\mu}_\tau > \mu_\tau$ iff

$$\begin{aligned} 1 - \frac{\lambda(1 - e^{-\psi\tau})}{1 - \mu(1 - e^{-\psi\tau})(1 - \lambda)} &> e^{-\lambda\psi\tau} \\ [1 - \mu(1 - e^{-\psi\tau})(1 - \lambda)](1 - e^{-\lambda\psi\tau}) &> (\lambda - 1)(1 - e^{-\psi\tau}) + (1 - e^{-\psi\tau}) \\ e^{-\psi\tau} - e^{-\lambda\psi\tau} &> (\lambda - 1)(1 - e^{-\psi\tau})(1 - \mu(1 - e^{-\lambda\psi\tau})) \\ \frac{e^{-\psi\tau} - e^{-\lambda\psi\tau}}{(1 - e^{-\psi\tau})(1 - \mu(1 - e^{-\lambda\psi\tau}))} &> \lambda - 1 \end{aligned}$$

As the LHS here tends to zero, $\psi > 0$, we have $\tilde{\mu}_\tau > \mu_\tau$ for τ large only if $\lambda < 1$. Thus iterated application of the bad news makes it more effective at depressing beliefs if $\lambda < 1$. But also notice the converse is true and if $\lambda > 1$ then one-off application of a lot of bad news depresses beliefs more than iteration.

Axiom on Order Reversal and its Implications

To describe this formally we consider a subset of signals: $A \subset S$. Then, we posit a pair of two-step procedures for learning the signal s : The first procedure begins with an experiment where the agent is told the signal s only if the realised s is in the set A . But if the realised signal s is not in the set A , then the agent is just told $s \notin A$. If they were told $s \notin A$ in the first experiment, then a second experiment is run where the agent learns the value of $s \in S \setminus A$.

The second procedure reverses this order. First the agent runs an experiment where they learn the value of s only if $s \in S \setminus A$. If the signal is not in this set then they are told just $s \in A$. If they were told $s \in A$ at the first stage, then a second experiment is run where they learn the value of $s \in A$. In Axiom 6 below we requires that both of these processes for learning the signal s result in the same terminal belief profile. We do not, however, require that either of these belief revision procedures has a terminal belief profile that is the same as the one-step process above.

Some additional notation is necessary for the formal statement of this Axiom. The set $A := \{1, 2, \dots, m\}$ and $S \setminus A := \{m + 1, \dots, n\}$. The fact that we restrict A to

only consist of the first m signals is unimportant, because of the symmetry axiom this will apply to all non-empty subsets of signals A . As A has m elements, the update after the first experiment of the first procedure is described by the function U_{m+1} . Whereas the update after the first experiment of the second procedure is described by the function U_{n-m+1} , because $S \setminus A$ has $n - m$ elements. In the first experiment of the first procedure the signals have the probabilities $p_A^\theta := (p_1^\theta, \dots, p_m^\theta, q_A^\theta) \in \Delta^o(m+1)$, where $q_A^\theta = 1 - \sum_{k=1}^m p_k^\theta$. Similarly in the first experiment in the second procedure the signals have the probabilities $p_{S \setminus A}^\theta = (q_{S \setminus A}^\theta, p_{m+1}^\theta, \dots, p_n^\theta) \in \Delta(n-m+1)$, where $q_{S \setminus A}^\theta = 1 - \sum_{k=m+1}^n p_k^\theta$. Thus the first experiment in the first procedure results in the profile of updated beliefs $U_{m+1}(\mu, (p_A^\theta)_{\theta \in \Theta})$ and the first experiment in the second procedure has the updated beliefs $U_{n-m+1}(\mu, (p_{S \setminus A}^\theta)_{\theta \in \Theta})$.

Axiom 6 (Order Independence). For all $\mathcal{E} = (S, (p^\theta)_{\theta \in \Theta})$, $\mu \in \Delta(\Theta)$, any $m \leq n$, and $n \geq 3$.

$$\begin{aligned} & \left[U_{m+1}^1(\mu, (p_A^\theta)_{\theta \in \Theta}), \dots, U_{m+1}^m(\mu, (p_A^\theta)_{\theta \in \Theta}), U_{n-m} \left(U_{m+1}^{m+1}(\mu, (p_A^\theta)_{\theta \in \Theta}), \left(\frac{\tilde{p}_{S \setminus A}^\theta}{q_A^\theta} \right)_{\theta \in \Theta} \right) \right] \\ & \equiv \\ & \left[U_m \left(U_{n-m+1}^1(\mu, (p_{S \setminus A}^\theta)_{\theta \in \Theta}), \left(\frac{\tilde{p}_A^\theta}{q_{S \setminus A}^\theta} \right)_{\theta \in \Theta} \right), U_{n-m+1}^2(\mu, (p_{S \setminus A}^\theta)_{\theta \in \Theta}), \dots, U_{n-m+1}^{n-m+1}(\mu, (p_{S \setminus A}^\theta)_{\theta \in \Theta}) \right] \end{aligned}$$

Where: $\tilde{p}_A^\theta := (p_1^\theta, \dots, p_m^\theta)$, $\tilde{p}_{S \setminus A}^\theta := (p_{m+1}^\theta, \dots, p_n^\theta)$.

Axiom 6 does not impose the condition that either of the terms in this equality is equal to the profile of beliefs in the one-off belief revision process: $U_n(\mu, (p^\theta)_{\theta \in \Theta})$. Axiom 4 says that when $m = 1$ both of these terms are equal to $U_n(\mu, (p^\theta)_{\theta \in \Theta})$. Then, an iterated application of Axiom 4 will imply they are equal for all m . Thus Axiom 6 is implied by Axiom 4 and symmetry, but potentially this axiom is weaker than divisibility.

To see that this is not actually the case, and Axiom 6 implies Axiom 4, consider this condition when $m = 1$. In the first procedure, where $s = 1$ or $s \neq 1$ is learned first the terminal belief profile is:

$$\left[U_2^1(\mu, (p_1^\theta, 1 - p_1^\theta)_{\theta \in \Theta}), U_{n-1} \left(U_2^2(\mu, (p_1^\theta, 1 - p_1^\theta)_{\theta \in \Theta}), \left(\frac{p_{-1}^\theta}{1 - p_1^\theta} \right)_{\theta \in \Theta} \right) \right]$$

Or, recalling the definition (2),

$$\left[u(\mu, (p_1^\theta)_{\theta \in \Theta}), U_{n-1} \left(u(\mu, (1 - p_1^\theta)_{\theta \in \Theta}), \left(\frac{p_{-1}^\theta}{1 - p_1^\theta} \right)_{\theta \in \Theta} \right) \right].$$

However, (when $m = 1$) the terminal belief profile after the second procedure for learning the signal is $U_n(\mu, (p^\theta)_{\theta \in \Theta})$. This is because under this procedure at the first stage the agent learns s if $s \geq 2$ and otherwise learns that $s < 2$. But this implies that the first stage of the second procedure completely reveals s and there is no additional learning at the second stage. Thus when $m = 1$ Axiom 6 implies

$$U_n(\mu, (p^\theta)_{\theta \in \Theta}) = \left[u(\mu, (p_1^\theta)_{\theta \in \Theta}), U_{n-1} \left(u(\mu, (1 - p_1^\theta)_{\theta \in \Theta}), \left(\frac{p_{-1}^\theta}{1 - p_1^\theta} \right)_{\theta \in \Theta} \right) \right],$$

which is precisely the condition in Axiom 4.

Proof of Proposition 2

Proof. The updating satisfies the axioms required for Proposition 1 to hold, hence it is characterised by a homeomorphism $F : \Delta(\Theta) \rightarrow \Delta(\Theta)$. We can, therefore, define the stochastic process followed by the shadow beliefs, $\{\tilde{\mu}^t\}_{t=0}^\infty$, where $\tilde{\mu}^t := F(\mu^t)$. Applying the updating rule (8) in period t with the signal s^t to the initial beliefs μ^t , we have that the updated beliefs satisfy

$$F(\mu^{t+1}) = F(u(\mu^t, (p_{s^t}^\theta)_{\theta \in \Theta})) = \left(\frac{F_1(\mu^t)p_{s^t}^1}{\sum_{\theta \in \Theta} F_\theta(\mu^t)p_{s^t}^\theta}, \dots, \frac{F_K(\mu^t)p_{s^t}^K}{\sum_{\theta \in \Theta} F_\theta(\mu^t)p_{s^t}^\theta} \right).$$

Substitution for $\tilde{\mu}^t$ then gives

$$\tilde{\mu}^{t+1} = \left(\frac{\tilde{\mu}_1^t p_{s^t}^1}{\sum_{\theta \in \Theta} \tilde{\mu}_\theta^t p_{s^t}^\theta}, \dots, \frac{\tilde{\mu}_K^t p_{s^t}^K}{\sum_{\theta \in \Theta} \tilde{\mu}_\theta^t p_{s^t}^\theta} \right).$$

Applying the above we can, then, write

$$\frac{\tilde{\mu}_\theta^t}{\tilde{\mu}_{\theta'}^t} = \frac{\tilde{\mu}_\theta^{t-1} p_{s^{t-1}}^\theta}{\tilde{\mu}_{\theta'}^{t-1} p_{s^{t-1}}^{\theta'}}.$$

This is the expression for the Bayesian updating of the shadow beliefs. Iterating this relation and taking logarithms we get

$$\frac{1}{t} \ln \frac{\tilde{\mu}_\theta^t}{\tilde{\mu}_{\theta'}^t} = \frac{1}{t} \ln \frac{\tilde{\mu}_\theta^0}{\tilde{\mu}_{\theta'}^0} + \frac{1}{t} \sum_{\tau=0}^{t-1} \ln \frac{p_{s^\tau}^\theta}{p_{s^\tau}^{\theta'}}, \quad \tilde{\mu}_{\theta'}^0 > 0.$$

When the parameter is θ , the terms in the above summation are independently and identically distributed with the expectation $H(p^\theta \| p^{\theta'}) = \sum_s p_s^\theta \ln \frac{p_s^\theta}{p_s^{\theta'}} \geq 0$. This is the relative entropy of the measures p^θ and $p^{\theta'}$. By the Strong Law of Large Numbers (Kallenberg (2002) p.73)

$$\lim_{t \rightarrow \infty} \frac{1}{t} \ln \frac{\tilde{\mu}_\theta^t}{\tilde{\mu}_{\theta'}^t} \rightarrow H(p^\theta \| p^{\theta'}), \quad \forall \theta' \neq \theta;$$

\mathbb{P}^θ almost surely. For θ' with $H(p^\theta \| p^{\theta'}) > 0$ this implies $\tilde{\mu}_{\theta'}^t \rightarrow 0$ \mathbb{P}^θ almost surely. And $\frac{\tilde{\mu}_\theta^t}{\tilde{\mu}_{\theta'}^t} = \frac{\tilde{\mu}_\theta^0}{\tilde{\mu}_{\theta'}^0}$ for θ' with $H(p^\theta \| p^{\theta'}) = 0$. Thus $\tilde{\mu}^t$ converges almost surely to $\tilde{\mu}^\infty \in \Delta(\Theta)$. As F is a homeomorphism this implies $\mu^t = F^{-1}(\tilde{\mu}^t)$ converges almost surely to $\mu^\infty = F^{-1}(\tilde{\mu}^\infty)$.

When the assumption $p^\theta \neq p^{\theta'}$ for all θ' holds then $H(p^\theta \| p^{\theta'}) > 0$ for all θ' and $\tilde{\mu}^\infty = e_\theta$. By (11) this implies that $\mu^\infty = F^{-1}(\tilde{\mu}^\infty) = e_\theta$. \square

Proof of Proposition 3

Proof. Applying what we know about the updating relation gives the following expression for the updated beliefs after the signal s :

$$\begin{aligned} & u(\mu, (p_s^\theta), (p_s^{\theta'})) \\ &= \left(f^{-1} \left(\frac{f(\mu_\theta) p_s^\theta}{f(\mu_\theta) p_s^\theta + (1 - f(\mu_\theta)) p_s^{\theta'}} \right), 1 - f^{-1} \left(\frac{f(\mu_\theta) p_s^\theta}{f(\mu_\theta) p_s^\theta + (1 - f(\mu_\theta)) p_s^{\theta'}} \right) \right). \end{aligned}$$

Thus if $(u_\theta, u_{\theta'}) = u(\mu, (p_s^\theta), (p_s^{\theta'}))$ are the updated beliefs in the two parameters conditional on the signal s , then

$$f(u_\theta) := \frac{f(\mu_\theta)p_s^\theta}{f(\mu_\theta)p_s^\theta + (1 - f(\mu_\theta))p_s^{\theta'}}.$$

This is also the shadow posterior on θ' . Now we calculate the expected posterior odds ratio, conditional on the parameter θ , that is

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

Hence, adding unity to the extremes of this equality gives

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

Now we will apply Jensens inequality. If $\frac{1}{f(\cdot)}$ is locally convex then we have

$$\frac{1}{f(E^\theta(u_\theta))} \leq E^\theta \left(\frac{1}{f(u_\theta)} \right) = \frac{1}{f(\mu_\theta)} \quad \Leftrightarrow \quad f(\mu_\theta) \leq f(E^\theta(u_\theta)).$$

When $f(\cdot)$ is increasing this implies $\mu_\theta \leq E^\theta(u_\theta)$ and the individual's belief in θ on average increases when it sees evidence consistent with θ .

However, when $\frac{1}{f(\cdot)}$ is locally concave then from Jensen's inequality we have

$$\frac{1}{f(E^\theta(u_\theta))} \geq E^\theta \left(\frac{1}{f(u_\theta)} \right) = \frac{1}{f(\mu_\theta)} \quad \Leftrightarrow \quad f(\mu_\theta) \geq f(E^\theta(u_\theta)).$$

When $f(\cdot)$ is increasing this implies $\mu_\theta \geq E^\theta(u_\theta)$ and the individual's belief in θ on average decreases when it sees evidence consistent with θ , thus the individual is biased against the truth.

Finally, we must show that $\frac{1}{f(\cdot)}$ cannot be concave on the entire interval $(0, 1)$. Suppose it were concave, then for any $x \in (0, 1)$ and $\varepsilon < x$

$$\frac{1}{f(x)} \geq \frac{1-x}{1+x-2\varepsilon} \frac{1}{f(\varepsilon)} + \frac{2x-2\varepsilon}{1+x-2\varepsilon} \frac{1}{f(\frac{1}{2}(1+x))},$$

as

$$x = \frac{1-x}{1+x-2\varepsilon} \varepsilon + \frac{2x-2\varepsilon}{1+x-2\varepsilon} \frac{1}{2}(1+x).$$

But as $\varepsilon \rightarrow 0$ the RHS of the first of these inequalities converges to infinity (as $f(\varepsilon) \rightarrow 0$). Thus $f(x) = 0$ for all $x \in (0, 1)$ a contradiction. \square

Proof of Proposition 4

Proof. We begin by considering the map induced by f from the odds ratio $\frac{\mu}{1-\mu}$ to the odds ratio $\frac{f(\mu)}{1-f(\mu)}$. Hence we define the function

$$\phi(r) := \frac{f\left(\frac{r}{1+r}\right)}{1 - f\left(\frac{r}{1+r}\right)}.$$

Observe that $\phi : \frac{\mu}{1-\mu} \mapsto \frac{f(\mu)}{1-f(\mu)}$ is differentiable and is one to one. Differentiation gives

$$\phi'(r) = \frac{1}{\left(1 - f\left(\frac{r}{1+r}\right)\right)^2} \frac{f'\left(\frac{r}{1+r}\right)}{(1+r)^2} = \left(\frac{1-\mu}{1-f(\mu)}\right)^2 f'(\mu),$$

where the final term comes from a substitution $\mu = \frac{r}{1+r}$. This gives the bounds $\beta^{-2} \min_{\mu \in R(\mu_\theta)} f'(\mu) \leq \phi'(r) \leq \beta^2 \max_{\mu \in R(\mu_\theta)} f'(\mu)$.

The updated odds ratio, $\frac{u_\theta(\mu, p_s)}{1-u_\theta(\mu, p_s)}$, is mapped to the shadow belief space by ϕ . So, we have

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

Or

$$(20) \quad \frac{u_\theta(\mu, p_s)}{1-u_\theta(\mu, p_s)} = \phi^{-1}\left(\phi\left(\frac{\mu_\theta}{1-\mu_\theta}\right) \frac{p_s^\theta}{p_s^{\theta'}}\right).$$

Now let us calculate the variance of the updated beliefs. This is

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

Here $\pi_s = \sum_\theta \mu^\theta p_s^\theta$ is defined to be the probability of signal s . A substitution from (20) then gives

$$\text{Var}\left[\frac{u_\theta(\mu, p_s)}{1-u_\theta(\mu, p_s)}\right] = \frac{1}{2} \sum_{s, s' \in S} \pi_s \pi_{s'} \left(\phi^{-1}\left(\phi\left(\frac{\mu^\theta}{1-\mu^\theta}\right) \frac{p_s^\theta}{p_s^{\theta'}}\right) - \phi^{-1}\left(\phi\left(\frac{\mu^\theta}{1-\mu^\theta}\right) \frac{p_{s'}^\theta}{p_{s'}^{\theta'}}\right)\right)^2.$$

Now we can apply the intermediate value theorem to this difference to get

$$\phi^{-1}\left(\phi\left(\frac{\mu^\theta}{1-\mu^\theta}\right) \frac{p_s^\theta}{p_s^{\theta'}}\right) - \phi^{-1}\left(\phi\left(\frac{\mu^\theta}{1-\mu^\theta}\right) \frac{p_{s'}^\theta}{p_{s'}^{\theta'}}\right) = \frac{1}{\phi'(x)} \left(\frac{p_s^\theta}{p_s^{\theta'}} - \frac{p_{s'}^\theta}{p_{s'}^{\theta'}}\right) \phi\left(\frac{\mu^\theta}{1-\mu^\theta}\right)$$

for some $\frac{x}{1+x} \in R(\mu_\theta)$. We have an upper bound $B := \beta^2 \max_{\mu \in R(\mu_\theta)} f'(\mu)$ on $\phi'(x)$. If this upper bound is substituted into the above expression for the variance we get

$$\begin{aligned} \text{Var}\left[\frac{u_\theta(\mu, p_s)}{1-u_\theta(\mu, p_s)}\right] &\geq B^{-2} \phi\left(\frac{\mu^\theta}{1-\mu^\theta}\right)^2 \frac{1}{2} \sum_{s, s' \in S} \pi_s \pi_{s'} \left(\frac{p_s^\theta}{p_s^{\theta'}} - \frac{p_{s'}^\theta}{p_{s'}^{\theta'}}\right)^2 \\ &= B^{-2} \phi\left(\frac{\mu^\theta}{1-\mu^\theta}\right)^2 \text{Var}\left[\frac{p^\theta}{p^{\theta'}}\right]. \end{aligned}$$

Hence a sufficient condition for the updating to satisfy (4) the condition for overreaction is

$$B^{-2} \phi\left(\frac{\mu^\theta}{1-\mu^\theta}\right)^2 \text{Var}\left[\frac{p^\theta}{p^{\theta'}}\right] \geq \left(\frac{\mu^\theta}{1-\mu^\theta}\right)^2 \text{Var}\left[\frac{p^\theta}{p^{\theta'}}\right].$$

Recalling the definition of $\phi(\cdot)$, this condition can be written more simply as

$$\frac{f(\mu^\theta)}{\mu^\theta} \frac{1-\mu^\theta}{1-f(\mu^\theta)} \geq B = \beta^2 \max_{\mu \in R(\mu_\theta)} f'(\mu).$$

A sufficient condition for this is $\beta^{-4} \geq \max_{\mu \in R(\mu_\theta)} f'(\mu)$, which is exactly the condition given in the Proposition. The application of the other bound on ϕ' gives the sufficient condition for under-reaction. \square