

# Modeling the Change of Paradigm: Non-Bayesian Reactions to Unexpected News<sup>☆</sup>

Pietro Ortoleva

*Division of the Humanities and Social Sciences, California Institute of Technology*

PRELIMINARY AND INCOMPLETE: PLEASE DO NOT CIRCULATE

---

## Abstract

Despite its normative appeal and widespread use, Bayes' rule has two well-known limitations: first, it does not predict how agents should react to an information to which they assigned probability zero; second, a sizable empirical evidence documents how agents systematically deviate from its prescriptions by overreacting to information that they deemed possible but "unexpected." By replacing Dynamic Consistency with a novel axiom, Dynamic Coherence, we characterize an alternative updating rule that is not subject to these limitations, but at the same time coincides with Bayes' rule for "normal" events. In particular, we model an agent with a utility function over consequences, a prior over priors  $\rho$ , and a threshold. In the first period she chooses the prior that maximizes the prior over priors  $\rho$  - a *la maximum likelihood*. As new information is revealed: if the chosen prior assigns to this information a probability above the threshold, she follows Bayes' rule and updates it. Otherwise, she goes back to her prior over priors  $\rho$ , updates it using Bayes' rule, and then chooses the new prior that maximizes the updated  $\rho$ . We also extend our analysis to the case of ambiguity aversion.

*Key words:* Bayes' Rule, Updating, Dynamic Consistency, Ambiguity Aversion.

*JEL classification:* D81, C61.

---

## 1. Introduction

One of the most widespread assumptions in economics is that agents update their beliefs using Bayes' rule. There seem to be many reasons for this popularity. First of all, it is an extremely intuitive procedure, so intuitive that it is often considered one of the features of rationality. Moreover, there are many situations in which Bayes' rule provides an accurate description of

---

<sup>☆</sup>I would like to thank Kim Border, Paolo Ghirardato, Federico Echenique, Leet Yariv, Leonardo Pejsachowicz, Gil Riella, and the participants at seminars at ASU, Caltech, and Collegio Carlo Alberto for useful comments and discussions.

the agents' reactions to information. In addition, its extremely convenient functional form facilitates its application in economic models. Despite these advantages, however, Bayes' rule has two well-known limitations. First, it has no prescription about how should the agent react to an information to which she assigned probability zero – Bayes' rule is simply not defined in that case. This limitation has well-known consequences: for example, the notion of Bayesian Nash Equilibrium is often criticized (and then refined) since it posits no restrictions on the beliefs of the agents out of equilibrium path – which is an immediate consequence of the fact that Bayes' rule has no prescription about how should these beliefs be formed.

Moreover, even in the cases in which Bayes' rule does apply, a sizable amount of evidence collected in the past two decades has documented that decision makers systematically depart from its prescriptions.<sup>1</sup> While these departures take many forms, one seems to be of particular relevance: people tend to violate Bayes' rule when they receive news that they did not foresee, i.e. information to which they originally assigned a small probability: “in violation of Bayes' rule most people tend to overreact to unexpected and dramatic news events” (De Bondt and Thaler (1985, pg. 804)). This tendency seems to persist even when appropriate incentives are given to the subjects, and also when subjects are experts on the area in which they have to make predictions.

The goal of this paper is to develop axiomatically an alternative updating rule that tries to maintain the elements of appeal of Bayes' rule, including its simplicity, but that also reconciles with the two limitations mentioned above. In particular, we model an agent who reacts according to Bayes' rule when “normal” news is given to her, but who might overreact, and change her prior beyond the bayesian prescriptions, when she receives some information that she “did not expect.” To give an example of the behavior we have in mind, think about an investor who is constructing a portfolio to allocate her wealth. To guide her decision, our investor has a belief over the returns of each possible investment, a belief that could originate from some properly calibrated economic model. As time goes by, our investor receives new information, e.g. from observing the stock market, and revises her belief. In normal times it is reasonable to expect that she revises her belief following Bayes' rule. But what happens if a big, shocking news is revealed, like a financial crisis? While such crisis might have been considered possible by our investor – in the sense that she might have assigned it a positive probability – it might also have been very *unlikely* for her – she could have assigned to it a very *small* probability. How will she react? The point is, in light of this unexpected event our investor might go beyond updating her prior using Bayes' rule: she might think that she used the wrong prior to begin with, given that it assigned such a small probability to the realized event, and that she should therefore pick a whole new prior based on the new information. The arrival of unexpected news might therefore lead her to a “change of paradigm” that entails a change of belief beyond Bays' rule. For example, if she was using an economic model to form her original belief, and this model assigned a small probability to the information that was later revealed, then she might question the validity of this model and look for a new one. In this paper we aim to characterize this behavior.

---

<sup>1</sup>For experimental evidence of violations of Bayes' rule in general, and for low probability events in particular, see, among others, Tversky and Kahneman (1974, 1981), Kahneman and Tversky (1982), Griffin and Tversky (1992), Grether (1992) Holt and Smith (2009) and the surveys in Camerer (1995), Rabin (1998).

Many generalizations of Bayes' rule have been proposed in the literature. In the non-axiomatic literature see for example the Jeffrey's rule, Mullainathan (2002), Rabin (2002), Mullainathan, Schwartzstein, and Shleifer (2008), Gennaioli and Shleifer (2009).<sup>2</sup> All of these papers, however, propose significant generalization of Bayes' law, instead of focusing on the violations of it for unexpected events, which is the goal of our paper. Moreover, none of these papers have an axiomatic foundation, making it hard to test which of them provides the most accurate description of the behavior. Within the surprisingly small axiomatic decision theory literature, Epstein (2006) and Epstein, Noor, and Sandroni (2008) generalize Bayes' rule by allowing overweighting or underweighting of evidence. (Consequences of Non-bayesian updating are discussed in Epstein, Noor, and Sandroni (2009).) These papers as well, however, present much more general models, and do not focus on the kind of violations we are interested in. Furthermore, they look at a framework of preferences over menus while we look at a standard dynamic Anscombe-Aumann setup. The recent work of Kochov (2009) studies a Non-Bayesian behavior in the presence of unforeseen contingencies – agents who fail to properly account for event that will take place in the non-immediate future, a behavior very different from the one we are interested in. Both the *lexicographic beliefs* of Blume, Brandenburger, and Dekel (1991), and the *conditional probability systems* of Myerson (1986a,b) address the issue of beliefs for null events, with important results on the agent's reaction to them. Neither of these theories, however, seem to extend to non-bayesian reactions to unlikely but non-null events, one of the goals of this paper; furthermore, they are based on preferences that depart from standard expected utility maximization even for the static case – for example in Blume, Brandenburger, and Dekel (1991) preferences are not fully Archimedean. By contrast, we wish to study an agent that is completely standard in the static case, but might have a non-standard reaction to news.

In the game theory literature Foster and Young (2003) present a model in which agents use a procedure similar to ours to learn which strategies are used by their opponents. In particular, just like in our model, their agents “test” whether their current belief is correct, and perform a paradigm change if this hypothesis is rejected. (They also call their model hypothesis testing.) Their analysis, however, is focused on learning in games, and does not apply to the standard framework. Also, it contains no prescription about how a new belief should be chosen if the current one is rejected, while this is an essential component of our findings. Finally, as opposed to what happens in our model, their agents do not act as standard Bayesian agents if an hypothesis is not rejected. In this sense, their model has a similar spirit in terms of when a paradigm change should take place in a game, but is very different in all other aspects.

Finally, in the literature on ambiguity a large attention is devoted to the issue of updating ambiguous beliefs, and since the Bayesian postulates seems problematic in that framework, many generalizations of them have been proposed.<sup>3</sup> These generalizations, however, are aimed

---

<sup>2</sup>See also Brav and Heaton (2002) and Brandt et al. (2004) for additional references to the behavioral literature.

<sup>3</sup>See, among others, Gilboa and Schmeidler (1993), Epstein and Le Breton (1993), Epstein and Schneider (2003), Maccheroni, Marinacci, and Rustichini (2006), Siniscalchi (2006), Hanany and Klibanoff (2006), and Ghirardato, Maccheroni, and Marinacci (2008).

at reconciling the standard approach with the presence of ambiguity aversion, and most of them reduce to Bayes' rule in the case in which the agent is an expected utility maximizer (ambiguity neutrality). By contrast, our main goal is to study violations of Bayes' rule even in the case of standard expected utility maximizers.<sup>4</sup>

We consider a standard dynamic Anscombe-Aumann model in which we observe the preferences of the agent before and after she receives some information about the state of the world. The main result of this paper is to provide an axiomatic foundation for what we call the Hypothesis Testing representation. According to this representation, the agent has a utility function  $u$  over consequences, a prior over priors  $\rho$ , and a threshold  $\epsilon$  between 0 and 1. In the first period, our agent chooses the prior  $\pi$  that maximizes her prior over priors  $\rho$ , and she uses it to form her preferences as an expected utility maximizer. In our example of the investor, we can see  $\rho$  as a belief over the possible economic models she can use (each of which entails a certain belief over the state of the world), where the agent uses the model that she considers the most likely – in a *maximum likelihood* fashion. As new information  $i$  is revealed, our agent acts as follows. If the probability that her prior assigned to the information is *above* the threshold  $\epsilon$ , i.e.  $\pi(i) > \epsilon$ , then the model is not rejected and she simply updates her prior  $\pi$  using Bayes' rule, thus acting like a standard agent. If, however, the probability that her belief assigned to the information is *below* the threshold, i.e.  $\pi(i) \leq \epsilon$ , then the model is rejected and our agent: goes back to her prior over priors  $\rho$ ; updates it using the additional information that she has received; then chooses the prior  $\pi'$  that maximizes her *updated* prior over priors; using this prior she forms her preferences maximizing the expected utility. That is, if the model is rejected by the data she goes back and picks the new maximum likelihood model, using the updated prior over priors.

The axiomatic foundation of the Hypothesis Testing representation maintains the standard axioms of Bayesian Updating in an Anscombe-Aumann setup, but replaces Dynamic Consistency with a novel axiom, Dynamic Coherence. The basic idea of this axiom is to impose that the agent behaves coherently also when she acts in a dynamically-inconsistent manner, or when she faces some information to which she assigned probability zero (null events). In this sense, Dynamic Coherence is neither stronger nor weaker than Dynamic Consistency: it is not stronger since it allows the agent to act in a dynamically-inconsistent manner, albeit in a regulated way; and it is not weaker than Dynamic Consistency since it restricts the behavior of the agent when she faces an information to which she assigned probability zero – events on which Dynamic Consistency has no bite.

We then extend our analysis to the case in which the agent is ambiguity averse, instead of a standard expected utility maximizer like we have assumed in the first part of the paper. Our starting point is one dynamic version of the model of Gilboa and Schmeidler (1989): every period the agent has not one, but a set of priors, and she evaluates each acts using the most pessimistic of them; when she receives new information, she updates every prior in the set using Bayes' rule. (This model has been characterized in Ghirardato, Maccheroni, and Marinacci (2008).) We then depart from this model in a way similar to how we departed

---

<sup>4</sup>We will address the case of ambiguity aversion as well (see Section 3), but our model for updating ambiguous beliefs will not reduce to Bayes' rule in the case of ambiguity neutrality.

from Bayesian updating. In the Hypothesis Testing representation with Ambiguity Aversion the agent has a prior over sets of priors  $\rho$  and a threshold  $\epsilon$ . In the first period she chooses the set of priors  $\Pi$  that maximizes the prior over sets of priors  $\rho$ . As new information is revealed, if the most pessimistic prior in  $\Pi$  assigns to the information a probability above the threshold  $\epsilon$ , then the agent updates every prior in  $\Pi$  following Bayes' rule. Otherwise, she updates the prior over sets of priors  $\rho$ , and chooses the new set of priors  $\Pi'$  that maximizes it.

The remainder of the paper is organized as follows. Section 2 presents the setup, the axiomatic foundation, the main representation. Section 3 introduces the extensions of the model to the case of ambiguity aversion. Section 4 concludes. The proofs appear in the appendix.

## 2. The Hypothesis Testing Model

### 2.1 Setup and foundations

We adopt a standard dynamic version of the Anscombe-Aumann setup. We have a finite (non-empty) set  $\Omega$  of states of the world, a  $\sigma$ -algebra  $\Sigma$  over  $\Omega$ , and a (non-empty) set  $X$  of consequences, which is assumed to be a compact subset of a metric space. Let  $\Delta(X)$  stand for the set of all Borel probability measures (lotteries) on  $X$ . Denote by  $\mathcal{F}$  the set of all simple acts, that is, the set of all finite-valued  $\Sigma$ -measurable functions  $f : \Omega \rightarrow \Delta(X)$ . With a standard abuse of notation, for any  $p \in \Delta(X)$  denote by  $p \in \mathcal{F}$  the constant act that yields the consequence  $p$  at every state  $\omega \in \Omega$ . For any  $A \in \Sigma$ ,  $f, g \in \mathcal{F}$ , denote by  $fAg \in \mathcal{F}$  the act that coincides with  $f$  in  $A$  and with  $g$  outside of it, that is,  $fAg(\omega) = f(\omega)$  for every  $\omega \in A$ , and  $fAg(\omega) = g(\omega)$  for every  $\omega \in \Omega \setminus A$ . For any function  $u : X \rightarrow \mathbb{R}$  and  $p \in \Delta(X)$ , denote by  $\mathbb{E}_p(u)$  the expected value of  $u$  with respect to  $p$ . As standard, metrize  $\mathcal{F}$  by the product Prokhorov metric, and define null states as follows.

**Definition 1.** For any preference relation  $\succeq$  on  $\mathcal{F}$ , we say that  $B \in \Sigma$  is  $\succeq$ -**null** if  $fBg \sim g$  for any  $f, g \in \mathcal{F}$ .

The primitive of our analysis is a class of non-degenerate preference relations  $\{\succeq_A\}_{A \in \Sigma}$ , where by  $\succeq_A$  we understand the preference of the agent after she receives the information  $A \in \Sigma$ , while we denote by  $\succeq = \succeq_\Omega$  the preference at time 0, before the agent receives any information.

We start by imposing two standard postulates.

**A.1 (Well-Behaved Standard Preferences (WbP)).** For any  $A \in \Sigma$ ,  $f, g, h \in \mathcal{F}$ :

1. (Continuity): the sets  $\{f' \in \mathcal{F} : f' \succeq_A f\}$  and  $\{f' \in \mathcal{F} : f \succeq_A f'\}$  are closed;
2. (Independence): for any  $\alpha \in (0, 1)$  we have

$$f \succeq_A g \Leftrightarrow \alpha f + (1 - \alpha)h \succeq_A \alpha g + (1 - \alpha)h;$$

3. (Monotonicity): if  $f(\omega) \succeq_A g(\omega)$  for all  $\omega \in \Omega$ , then  $f \succeq_A g$ .
4. (Constant Preference Invariance): for any  $p, q \in \Delta(X)$ ,  $p \succeq_A q \Leftrightarrow p \succeq q$

**A.2 (Consequentialism (C)).** For any  $A \in \Sigma$ , and  $f, g \in \mathcal{F}$ , if  $f(\omega) = g(\omega)$  for all  $\omega \in A$ , then  $f \sim_A g$ .

Axiom WbP (Axiom 1) is a collection standard postulates that guarantee that both before and after the arrival of information our agent acts like a standard expected utility maximizer (part (1), (2) and (3)), and that the arrival of information does not affect the agent's ranking of the consequences in  $\Delta(X)$  (part (4)). Consequentialism (Axiom 2) is another standard postulate that guarantees that the agent believes in the information she receives: if she is told that the true state lies inside some  $A \in \Sigma$ , then she is indifferent between two acts that differ only outside of  $A$ .

We now turn to restrict the way beliefs evolve with information. To this end, the standard postulate is Dynamic Consistency.

**A.3 (Dynamic Consistency (DC)).** For any  $A \in \Sigma$ ,  $A$  not  $\succeq$ -null, and for any  $f, g \in \mathcal{F}$ , we have

$$f \succeq_A g \Leftrightarrow fAg \succeq g.$$

The basic idea of Dynamic Consistency is that the arrival of some information  $A \in \Sigma$  should not modify the ranking of two acts that coincide outside of  $A$ . It is well known that adding Dynamic consistency to WbP and Consequentialism implies that the agent follows Bayes' rule. (We refer to Ghirardato (2002) for an in-depth discussion of this standard postulate of its implications.) While appealing, however, Dynamic Consistency has two important limitations. First of all, it disciplines agent's preferences only if the event  $A$  is not  $\succeq$ -null, that is only if  $A$  is assigned a positive probability to begin with. This implies that any theory that derives from Dynamic Consistency is bound to have no predicting power on the agent's reaction to events to which she assigns probability zero. Second, as we mentioned in the introduction, a sizable experimental evidence documents systematic violations of Bayes' rule in general, and of Dynamic Consistency in particular. This seems to be especially true when the revealed information was assigned a small probability. In other words, despite its normative appeal on the one hand Dynamic Consistency seems to be too strong to be accepted as a positive axiom, while on the other hand it seems not strong enough since it doesn't restrict the agent's behavior in the case of null events.

To reconcile with the empirical evidence, we then replace Dynamic Consistency with an axiom that allows the agent to behave in a dynamically-inconsistent manner, but that guarantees that she does so in a regulated manner. At the same time, we would like such regularity to apply also to the agent's reaction to null events, so that we can develop a theory that applies to that case as well. The basic idea is that want to rule out the possibility that our agent could have a *circular* reaction to information. To illustrate, consider three possible events  $A_1, A_2, A_3 \in \Sigma$  such that: after being told that the true state lies either in  $A_1$  or in  $A_2$ ,

then the agent is *sure* that it lies in  $A_2$ ; if instead she were told that it lies either in  $A_2$  or in  $A_3$ , than she is sure that it lies in  $A_3$ ; but if she is told that it lies either in  $A_3$  or in  $A_1$ , then she is *sure* that it lies in  $A_1$ . That is, being told that the true state could be in one event or the next, she thinks it is in the latter; but were she told that it's either in the last or in the first, then she is sure it is in the first. Now, this would be naturally true if the agent had the *same* preferences after each of these events, i.e. if  $\succeq_{A_1} = \succeq_{A_2} = \succeq_{A_3}$ : if this were the case, then after each of these events she would be sure that the true state lies in  $A_1 \cap A_2 \cap A_3$ , giving us the beliefs described above. But, what if this is not the case, i.e.  $\succeq_{A_1} \neq \succeq_{A_2} \neq \succeq_{A_3}$ ? Then the agent's beliefs react to information in a *circular* manner: beliefs keep changing, but they form a loop, since at the end the agent is sure that the true state lies in the *first* event. Notice that this circularity could never take place when an agent is Bayesian, at least if  $A_1, A_2$ , and  $A_3$  are non-null: Bayes's rule would prevent it by construction. But of course Bayes' rule is stronger than simply postulating the lack of this circularity – it entails a much stronger form of consistency. By contrast, in the next axiom we simply posit that this circularity could not take place, for cycles of any length.

**A.4 (Dynamic Coherence).** *There are no  $A_1, \dots, A_n \in \Sigma$  such that  $\succeq_{A_1} \neq \succeq_{A_n}$ ,  $(\Omega \setminus A_{i+1})$  is  $\succeq_{(A_i \cup A_{i+1})}$ -null for  $i = 1, \dots, (n - 1)$ , and  $(\Omega \setminus A_1)$  is  $\succeq_{(A_n \cup A_1)}$ -null.*

It is not hard to see that Dynamic Coherence is neither stronger nor weaker than Dynamic Consistency. It is not stronger since it does allow violations of Dynamic Consistency, albeit regulating them – as we mentioned before, Bayes' rule is stronger than the lack of circularity that we posit with Dynamic Coherence. At the same time, it is not weaker than Dynamic Consistency since it disciplines the reaction to null events, on which Dynamic Consistency has no bite.<sup>5</sup>

## 2.2 The model

We are now ready to introduce our main representation. For simplicity, let us define a notation for Bayesian Updating. For any  $\pi \in \Delta(\Omega)$  and  $A \in \Sigma$  such that  $\pi(A) \geq 0$ , define  $\text{BU}(\pi, A) \in \Delta(\Omega)$  (bayesian update of  $\pi$  using  $A$ ) as

$$\text{BU}(\pi, A)(B) := \frac{\pi(A \cap B)}{\pi(A)}$$

for all  $B \in \Sigma$ . Abusing notation, for any  $\rho \in \Delta(\Delta(\Omega))$  and  $A \in \Sigma$  such that  $\pi(A) > 0$  for some  $\pi \in \text{supp}(\rho)$ , define also

$$\text{BU}(\rho, A)(\pi) := \frac{\pi(A)\rho(\pi)}{\int_{\Delta(\Omega)} \pi(A)\rho(\pi)d\pi}$$

---

<sup>5</sup>This implies that if there are no null events, then Dynamic Coherence is strictly weaker than Dynamic Consistency. On the other hand, if null events exist, then it would make sense to impose both Dynamic Coherence and Dynamic Consistency: the representation theorem in Section 2.3 will analyze this case as well.

for all  $\pi \in \Delta(\Omega)$ .

**Definition 2.** A class of preferences relations  $\{\succeq_A\}_{A \in \Sigma}$  admits an **Hypothesis Testing Representation** if there exists a continuous function  $u : X \rightarrow \mathbb{R}$ , a prior  $\rho \in \Delta(\Delta(\Omega))$  with finite support, and  $\epsilon \in [0, 1)$  such that for any  $A \in \Sigma$  there exist  $\pi_A \in \Delta(\Omega)$  such that:

1. for any  $f, g \in \mathcal{F}$

$$f \succeq_A g \Leftrightarrow \sum_{\omega \in \Omega} \pi_A(\omega) \mathbb{E}_{f(\omega)}(u) \geq \sum_{\omega \in \Omega} \pi_A(\omega) \mathbb{E}_{g(\omega)}(u);$$

2.  $\{\pi_\Omega\} = \arg \max_{\pi \in \Delta(\Omega)} \rho(\pi)$ ;

- 3.

$$\pi_A = \begin{cases} \text{BU}(\pi_\Omega, A) & \text{if } \pi_\Omega(A) > \epsilon \\ \text{BU}(\pi_A^*, A) & \text{otherwise} \end{cases}$$

where  $\{\pi_A^*\} = \arg \max_{\pi \in \Delta(\Omega)} \text{BU}(\rho, A)(\pi)$ ;

4. for any  $A \in \Sigma$  there exist  $\pi \in \text{supp}(\rho)$  such that  $\pi(A) > 0$ .

In an Hypothesis testing representation the agent has a utility function, a prior over priors  $\rho$ , and a threshold  $\epsilon$ . In the first period she chooses the prior  $\pi$  that maximizes the prior over priors  $\rho$ . She behaves as if she were choosing which “theory” to use to forecast the states of the world: given a certain belief  $\rho$  over the possible theories, she picks the most likely one – in a maximum likelihood fashion. As new information  $A$  is revealed, two things can happen. If the prior she was using assigned to this event a probability *above* the threshold, i.e.  $\pi(A) > \epsilon$ , then our agent “keeps” her prior, and simply updates it with Bayes’ rule. That is, if the information is not unexpected, then we are in the “business as usual” situation, and our agent behaves like a standard one. If, however, she is given an information that she did not expect, that is, if the likelihood that her prior assigned to that information is *below* the threshold, i.e.  $\pi(A) \leq \epsilon$ , then our agent revises which prior to use. It is as if she thought: “If the prior I am using did not forecast what happened, then, maybe, it is the wrong prior!” From this point of view, our agent acts as if she were “testing” her prior and “rejecting” the hypothesis that it is correct if it falls below the threshold, as if it were a confidence level – hence the name of the model. And how does she choose a new prior? She updates her prior over priors using Bayes’ rule; then she chooses the prior that maximizes the *updated* prior over priors; and finally she uses this prior until a new information is revealed.

The decision rule that we have just described in words, however, might leave some room to indeterminacy: what would happen if there is more than one prior that maximizes the updated prior over prior? Which one is chosen? Such indeterminacy could not take place in an Hypothesis Testing model, since  $\rho$  is constructed in such a way that the argmax of the updated prior over priors is always *unique*.

Just like Dynamic Coherence is neither stronger nor weaker than Dynamic Consistency, an Hypothesis Testing representation is neither more general nor more restrictive than the

standard Bayesian model. To see why, notice that if  $\epsilon = 0$ , then our agent behaves exactly like a standard Bayesian agent whenever Bayes' rule applies; but her behavior is disciplined also when Bayes' rule does not apply (null events), thus extending the predictive power of the theory to all possible events and making this model a special case of the Bayesian one. On the other hand, when  $\epsilon > 0$  the Hypothesis Testing representation allows non-bayesian reactions to non-null events. In this sense, the Hypothesis Testing representation generalizes the Bayesian approach to allow for that over-reaction to unexpected news that has been documented empirically.<sup>6</sup>

Since there could be multiple values of  $\epsilon$  that represent the same preferences (see Section 2.3), we focus on the representations with the smallest possible values of  $\epsilon$ .

**Definition 3.** An Hypothesis Testing Representation  $(u, \rho, \epsilon)$  is **minimal** if there is no  $\epsilon' \in [0, 1)$  such that  $\epsilon' < \epsilon$  and  $(u, \rho, \epsilon')$  is an Hypothesis Testing Representation of the same preferences.

### 2.3 Representation Theorem

**Theorem 1.** A class of preference relations  $\{\succeq_A\}_{A \in \Sigma}$  satisfies *WbP*, *Consequentialism*, and *Dynamic Coherence* if and only if it admits a minimal Hypothesis Testing Representation  $(u, \rho, \epsilon)$ .

Moreover,  $\epsilon = 0$  if and only if  $\{\succeq_A\}_{A \in \Sigma}$  satisfies also *Dynamic Consistency*.

Theorem 1 shows that by replacing Dynamic Consistency with Dynamic Coherence we obtain exactly the Hypothesis Testing Representation. Moreover, it shows that Dynamic Consistency and Dynamic Coherence together guarantee that our agent behaves like a standard Bayesian agent whenever Bayes' rule applies, but reconsider which prior to use whenever she faces an information to which she assigned probability zero.

We now turn to discuss the uniqueness properties of an Hypothesis Testing Representation. It is standard practice to show that the utility function is unique up to a positive affine transformation. The threshold  $\epsilon$  is unique, but only thanks to our focus on minimal representations: in general there might be a continuum of values of  $\epsilon$  that would work.<sup>7</sup> As for the prior over priors  $\rho$ , it turns out that it is not unique, and that even its support is not unique. There are essentially three reasons why this is the case. (To avoid confusion we address the elements of the support of  $\rho$  as “models.”) First, we can always add to the support of  $\rho$  an additional model with a likelihood so low that it will never be used. That

---

<sup>6</sup>More precisely, if  $\epsilon = 0$  then the Hypothesis Testing model is a special case of the Bayesian one. The converse is true if there are no null events. This means that if both conditions hold then the two models coincide.

<sup>7</sup>To see why, call  $a$  the likelihood of the *least* likely event that does not trigger a violation of Bayes' rule, and call  $b$  the likelihood of the *most* likely event that does trigger a violation of Bayes' rule. Then, we must have  $b < a$ , and any  $\epsilon \in [b, a)$  would work. The reason is, our space of events might not be “dense” in this sense, and therefore the threshold  $\epsilon$  is not uniquely identified. On the other hand, any  $\epsilon$  must lie in this range.

is, we can always add “redundant” models leaving the behavior unaffected. Second, even if we removed these redundant ones, we are bound to identify models only after the events that trigger the agent to use them, and therefore we have no control over what these models prescribe outside of these events.<sup>8</sup> For the same reason, moreover, there might be multiple ways to combine these models, making also the cardinality of the support of  $\rho$  not unique.<sup>9</sup>

#### 2.4 Discussion: the Hypothesis Testing model as a form of bounded rationality

One characteristic of the Hypothesis Testing model is that agents forms beliefs in a non-standard way even in the first period, before any information is revealed. In fact, we can think of the agent as if she had in mind a set of “conceivable” priors (the support of  $\rho$ ), which she can rank in terms of plausibility (forming  $\rho$ ), of which, however, she uses only one, the most likely one – a’ la maximum likelihood.<sup>10</sup> The other priors in the support of  $\rho$ , albeit “conceivable,” are in fact not used to make decisions unless some unexpected news is revealed.<sup>11</sup> This is in contrast to the behavior prescribed by the standard Bayesian approach, according to which if the agent has a prior over priors, she should choose as a belief its *expectation*, not its maximizer. That is, she should consider all conceivable priors and weight them appropriately, instead of using only the most likely one. While these two approaches are behaviorally indistinguishable in a static case, since all we can see is the belief that is used and not where it comes from, it is the analysis of the dynamic case that allow us to set them apart. In particular, it is by observing violations of Dynamic Consistency that we can infer how has the agent formed her beliefs in the first place.

While seemingly irrational, we consider the use of only a *subset* of all possible theories as rather realistic. To see why, let us go back to our example of the investor. When she chooses a portfolio, our investor could indeed considers *all* possible economic and statistical models that she can come up with, assign to them a relative likelihood, estimate them all, look at their

---

<sup>8</sup>Consider a model  $\pi$  used after the arrival of some information  $A \in \Sigma$ . Since all we observe is  $\text{BU}(\pi, A)$ , then we can replace  $\pi$  with any  $\pi'$  such that  $\text{BU}(\pi|A) = \text{BU}(\pi'|A)$  and  $\pi(A) = \pi'(A)$ .

<sup>9</sup>Consider four events  $A, B, C, D \in \Sigma$ , all with a likelihood below  $\epsilon$ , such that  $A \supset B$ ,  $C \supset D$  and  $A \cap C = \emptyset$ . Denote by  $\pi_i$  the priors that are used after event  $i$ . Assume that we have  $\text{BU}(\pi_A, B) \neq \pi_B$  and  $\text{BU}(\pi_C, D) \neq \pi_D$ . There are at least three ways to represent this case. First, we can create four distinct models, one for each event, giving to  $\rho$  the appropriate weights. Second, we can have only two models,  $\pi_1$  and  $\pi_2$ , such that  $\text{BU}(\pi_1, A) = \pi_A$ ,  $\text{BU}(\pi_1, C) = \pi_C$ ,  $\text{BU}(\pi_2, B) = \pi_B$ ,  $\text{BU}(\pi_2, D) = \pi_D$ . (Model  $\pi_1$  is used after  $A$  and  $C$ , while model  $\pi_2$  is used after  $B$  and  $D$ .) Third, we could use two other models,  $\pi_3$  and  $\pi_4$ , such that  $\text{BU}(\pi_3, A) = \pi_A$ ,  $\text{BU}(\pi_3, D) = \pi_D$ ,  $\text{BU}(\pi_4, B) = \pi_B$ ,  $\text{BU}(\pi_4, C) = \pi_C$ . (Model  $\pi_3$  is used after  $A$  and  $D$ , while model  $\pi_4$  is used after  $B$  and  $C$ .) Unfortunately, this example also shows that we cannot find a unique  $\rho$  with “minimal” support.

<sup>10</sup>Of course this chosen prior might itself be the combination of multiple ones with given weight. For example, the many priors in the support of  $\rho$  might originate from different weights put on a fixed set of beliefs.

<sup>11</sup>This suggests that our agent might have two layers of belief, what is conceivable but not used, and what is actually used. The idea of multiple layers of beliefs is explored also in Blume, Brandenburger, and Dekel (1991), who consider the *lexicographic probabilities*. As we mentioned before, however, their model studies the reactions to null-events by allowing non-standard (non-archimedean) preferences in the static case. By contrast, our model has standard preferences on the static case but a non-standard dynamic, and moreover allows for non-standard reactions to events that are not null but simply low probability.

predictions, and finally come up with a belief which is the weighted average of the predictions of all of these models. This is a behavior prescribed by the standard approach. Alternatively, our agent could choose just a subset of these models, the ones that she considers the most likely, estimate only these, and then use as a belief the weighted average of the predictions of only this selected subset of models.<sup>12</sup> She would then consider alternative models only if some unexpected news is revealed. The latter behavior seems more realistic to us.

This leaves us with the question of why should the agent adopt this behavior, since by restricting her attention to a specific subset of models she could end up with inaccurate predictions. There are at least four reasons why we believe this could be the case. First, it could be a rational reaction to a form of bounded rationality/costly thinking/cost of considering models. In fact, if estimating a model is costly, as in most cases it is, then our agent has an incentive to be parsimonious in the number of models she considers.<sup>13</sup> Such cost could also be psychological: agents might not like to be reminded that they have a limited knowledge of the reality, and that they do not even know which model to use, and might therefore prefer to focus on a single one as if it were true – a tendency highly emphasized in the psychology literature. Second, it could be seen simply as another instance of the standard behavioral bias that leads the agents to disregard low probability events.<sup>14</sup> Third, this behavior could stem for a *preference for simplicity*: considering only a few models is simpler, and the agent might prefer it in an Occam’s Razor sense. Finally, it turns out that considering such simpler theories might actually be *optimal* even if there were no cost of considering models or other additional costs. In fact, Gilboa and Samuelson (2008) show that considering complex theories might induce people to *overfit* the data and engage in ineffective learning, generating worse predictions. By contrast, simpler theories do not have this problems, and might lead to optimal behavior. We refer to their work for an detailed analysis.

### 3. The Hypothesis-Testing model with ambiguity aversion

Our analysis thus far was carried out under the assumption that in every period, and after every information, the decision maker behaves like a standard expected utility maximizer with a well-formed (and unique) prior over the states of the world. This is an immediate consequence of axiom WbP (Axiom 1) which includes independence. Since the Ellsberg paradox, however, it is well known that independence is often violated, and agents do not behave like standard expected utility maximizers with a single prior over the states of the world. Rather, they are shown to be ambiguity averse: instead of using expected utility, they dislike betting on outcomes that depend on the realization of unknown states of the world, and have a preference for hedging. This seems all the more problematic in our analysis since

---

<sup>12</sup>Once again, notice that the chosen prior in an Hypothesis Testing representation could come from averaging a *set* of models with a relative weight.

<sup>13</sup>In fact, it is easy to see that we could have re-written our representation as a representation in which the agent is choosing the optimal number of models given a cost of considering each of them.

<sup>14</sup>That is, the models to which  $\rho$  assigns a small probability could be disregarded just like low probability events are often disregarded in the choice among gambles.

the instances in which we should expect a violation of the Bayesian model are also the ones in which we should expect ambiguity aversion. The goal of this section is to show that an appropriate extension of our results, using the same axiom, Dynamic Coherence, would hold true also in the more general case in which agents are ambiguity averse.

The issue of updating beliefs under ambiguity has been studied by an extensive literature.<sup>15</sup> Almost all of these works, however, have focused on the complicated interaction between ambiguity aversion and Bayesian updating, without questioning the updating procedure in the first place. In particular, most of these model reduce to the case of Bayes' rule if the agent is ambiguity neutral (expected utility maximizer). By contrast, in Section 2 we have suggested an alternative to Bayesian updating in the case of ambiguity neutrality, and we now wish to do the same in the case of ambiguity aversion.

### 3.1 Foundations with Ambiguity Aversion

In what follows we model ambiguity aversion using the well-known model of Gilboa and Schmeidler (1989). In particular, we replace standard Independence with their's C-independence and Ambiguity Aversion axioms, and obtain WbP-AA (Axiom 5) to replace WbP (Axiom 1). We refer to Gilboa and Schmeidler (1989) for a more detailed discussion.

#### A.5 (Well-Behaved Standard Preferences with Ambiguity Aversion (Wbp-AA)).

For any  $A \in \Sigma$ ,  $f, g, h \in \mathcal{F}$ :

1. (Continuity): the sets  $\{f' \in \mathcal{F} : f' \succeq_A f\}$  and  $\{f' \in \mathcal{F} : f \succeq_A f'\}$  are closed;
2. (C-Independence): for any  $\alpha \in (0, 1)$ ,  $x \in \Delta(X)$

$$f \succeq_A g \Leftrightarrow \alpha f + (1 - \alpha)x \succeq_A \alpha g + (1 - \alpha)x;$$

3. (Uncertainty Aversion) for any  $\alpha \in (0, 1)$ , if  $f \sim_A g$  then  $\alpha f + (1 - \alpha)g \succeq_A f$ .
4. (Monotonicity): if  $f(\omega) \succeq_A g(\omega)$  for all  $\omega \in \Omega$ , then  $f \succeq_A g$ .
5. (Constant Preference Invariance): for any  $B \in \Sigma$ ,  $p, q \in \Delta(X)$ ,  $p \succeq_A q \Leftrightarrow p \succeq_B q$

It is well-known that WbP-AA (Axiom 5) guarantees that in every period and after every information the agent has a closed and convex set of prior beliefs over the states of the world (instead of a unique prior), and judges every act by the expected utility computed with the most pessimistic prior for that act. (Again, we refer to Gilboa and Schmeidler (1989) for a discussion on the properties of this representation.)

While we replace WbP (Axiom 1) with WbP-AA above (Axiom 5), we posit that Consequentialism and Dynamic Coherence hold as before also for the case with Ambiguity Aversion

---

<sup>15</sup>See, among others, Gilboa and Schmeidler (1993), Epstein and Le Breton (1993), Epstein and Schneider (2003), Maccheroni, Marinacci, and Rustichini (2006), Siniscalchi (2006), Hanany and Klibanoff (2006), and Ghirardato, Maccheroni, and Marinacci (2008).

– with the same intuition. This is in contrast with what happens with Dynamic Consistency, since it is well known that when agents are ambiguity averse, especially a’ la Gilboa and Schmeidler (1989), Dynamic Consistency might be too strong of a requirement. For example, consider an agent who updates her set of priors by updating each prior using Bayes’ rule: Epstein and Schneider (2003) show that such an agent might violate Dynamic Consistency.<sup>16</sup> In fact, Ghirardato et al. (2008) show that this model is equivalent to the agent satisfying a weakening of Dynamic Consistency, in which this axiom applies only to a *subset* of the original preference relation, which they call the “unambiguously preferred” relation. In particular, following the analysis in Ghirardato et al. (2004) and Ghirardato et al. (2008), for any  $A \in \Sigma$  define the preference relation  $\succeq_A^*$  as follows:

$$f \succeq_A^* g \quad \text{if} \quad \lambda f + (1 - \lambda)h \succeq_A \lambda g + (1 - \lambda)h$$

for all  $\lambda \in [0, 1]$  and all  $h \in \mathcal{F}$ . (It is not hard to see that  $\succeq_A^*$  is the largest restriction of  $\succeq_A$  that satisfies independence, and we clearly have  $\succeq_A^* = \succeq_A$  if the latter satisfies independence.) The idea of the axiom is to impose a dynamically consistent behavior on this restricted preference.

**A.6 (Restricted Dynamic Consistency (RDC)).** *For any  $A \in \Sigma$ ,  $A$  not  $\succeq$ -null, and for any  $f, g \in \mathcal{F}$ , we have*

$$f \succeq_A^* g \Leftrightarrow fAg \succeq^* g.$$

We refer to Ghirardato, Maccheroni, and Marinacci (2008) for further discussion. Notice that if every preference satisfies independence (Axiom 1), then Restricted Dynamic Consistency (Axiom 6) is clearly equivalent to standard Dynamic Consistency (Axiom 3).<sup>17</sup>

### 3.2 The representation with Ambiguity Aversion

In the case of ambiguity aversion modeled a’ la Gilboa and Schmeidler (1989) the agent has not one, but a set of priors. As she receives new information this set of priors will be modified: a natural extension of the standard Bayesian approach to this case is the model in which the agent reacts to this new information *by updating using Bayes’ rule each of the priors in her set of priors*. Just like in the first part of the paper we started from the Bayesian model and defined the Hypothesis Testing one, in our analysis here we start from this generalization to extend the Hypothesis Testing model to the case of ambiguity aversion. The idea of this generalization is that our agent has a prior over *sets* of priors, and picks the

---

<sup>16</sup>In particular, they show that the agent satisfies Dynamic Consistency if and only if the set of priors they use in the first period satisfies a property called *rectangularity*. But this means that a postulate that is supposed to regulate only the reaction to information is actually constraining the beliefs even before any information is revealed.

<sup>17</sup>Moreover, it turns out that if there are no null events, then under WbP and Consequentialism we have that Dynamic Coherence is weaker than Restricted Dynamic Consistency. (This is an immediate consequence of Theorem 5 below.)

set of priors that maximizes this prior. Then, she tests this set of prior using a threshold: if it passes the test, she updates every prior in the set using Bayes' rule. Otherwise, she updates her prior over sets of priors, and chooses the new set that maximizes it.

For simplicity, we introduce a notation for the extension of the notion of Bayesian Updating to updating a set of priors and a prior over sets of priors. By  $\mathcal{G}$  denote the set of closed and convex subsets of  $\Delta(\Omega)$ . For any  $\Pi \subseteq \mathcal{G}$  and  $A \in \Sigma$  such that  $\pi(A) > 0$  for some  $\pi \in \Pi$ , define  $\hat{B}U(\Pi, A)$  as

$$\hat{B}U(\Pi, A) := \{\pi \in \Delta(\Omega) : \pi = \text{BU}(\pi', A) \text{ for some } \pi' \in \Pi\}.$$

Moreover, for any  $\rho \in \Delta(\mathcal{G})$  and  $A \in \Sigma$  such that for some  $\Pi \in \text{supp}(\rho)$  we have  $\pi(A) > 0$  for all  $\pi \in \Pi$ , define  $\overline{B}U(\rho, A)$  as

$$\overline{B}U(\rho, A)(\Pi) := \frac{\min_{\pi \in \Pi} \pi(A) \rho(\Pi)}{\int_{\mathcal{G}} \min_{\pi \in \Pi} \pi(A) \rho(\Pi) d\Pi}.$$

**Definition 4.** A class of preferences relations  $\{\succeq_A\}_{A \in \Sigma}$  admits an **Hypothesis Testing Representation with Ambiguity Aversion** if there exists a function  $u : X \rightarrow \mathbb{R}$ , a prior  $\rho \in \Delta(\mathcal{G})$  with finite support, and  $\epsilon \in [0, 1]$  such that for any  $A \in \Sigma$  there exist a  $\Pi_A \in \mathcal{G}$  such that:

1. for any  $f, g \in \mathcal{F}$

$$f \succeq_A g \Leftrightarrow \min_{\pi \in \Pi_A} \sum_{\omega \in \Omega} \pi(\omega) \mathbb{E}_{f(\omega)}(u) \geq \min_{\pi \in \Pi_A} \sum_{\omega \in \Omega} \pi(\omega) \mathbb{E}_{g(\omega)}(u);$$

2.  $\{\Pi_\Omega\} = \arg \max_{\Pi \in \mathcal{G}} \rho(\Pi)$ ;

- 3.

$$\Pi_A = \begin{cases} \hat{B}U(\Pi_\Omega, A) & \text{if } \pi(A) > \epsilon \text{ for all } \pi \in \Pi_\Omega \\ \hat{B}U(\Pi_A^*, A) & \text{otherwise} \end{cases}$$

where  $\{\Pi_A^*\} = \arg \max_{\Pi \in \mathcal{G}} \overline{B}U(\rho, A)(\Pi)$ ;

4. for any  $A \in \Sigma$  there exist  $\Pi \in \text{supp}(\rho)$  such that  $\pi(A) > 0$  for all  $\pi \in \Pi$ .

**Definition 5.** An Hypothesis Testing Representation with Ambiguity Aversion  $(u, \rho, \epsilon)$  is **minimal** if there is no  $\epsilon' \in [0, 1)$  such that  $\epsilon' < \epsilon$  and  $(u, \rho, \epsilon')$  is an Hypothesis Testing Representation with Ambiguity Aversion of the same preferences.

Finally, the representation theorem.

**Theorem 2.** A class of preference relations  $\{\succeq_A\}_{A \in \Sigma}$  satisfies *Wbp-AA*, *Consequentialism*, and *Restricted Dynamic Coherence* if and only if it admits a minimal Hypothesis Testing Representation with Ambiguity Aversion  $(u, \rho, \epsilon)$ .

Moreover,  $\epsilon = 0$  if and only if  $\{\succeq_A\}_{A \in \Sigma}$  satisfies also *Restricted Dynamic Consistency*.

Theorem 2 shows that the same Dynamic Coherence that we have imposed for the standard case gives us the desired representation also for the case of Ambiguity Aversion. Moreover, when Restricted Dynamic Consistency is satisfied as well, we obtain a representation in which the agent updates the set of prior using Bayes' rule if she faces a non-null event, and picks a new set of prior otherwise.

#### 4. Conclusion

In this paper we have developed axiomatically an alternative to the standard Bayesian model. We study an agent who behaves like a standard Bayesian when she receives an information that is not "unexpected," i.e. to which she assigned a probability above a threshold. If this is not the case, however, instead of following Bayes' rule she reconsiders her prior by updating a prior over priors and picking the most likely one after the update. We have also discussed extensions of the model to the case in which the preferences are ambiguity averse.

#### Appendix: Proofs

*Proof of Theorem 1*

[*Sufficiency of the Axioms*] Given Axiom 1, it is standard practice to show that for any  $A \in \Sigma$ , there exist  $u_A : X \rightarrow \mathbb{R}$ ,  $\pi_A \in \Delta(X)$  such that for any  $f, g \in \mathcal{F}$

$$f \succeq_A g \Leftrightarrow \sum_{\omega \in \Omega} \pi_A(\omega) \mathbb{E}_{f(\omega)}(u_A) \geq \sum_{\omega \in \Omega} \pi_A(\omega) \mathbb{E}_{g(\omega)}(u_A), \quad (\text{A.1})$$

where  $\pi_A$  is unique and  $u_A$  is unique up to a positive affine transformation. It is also standard practice to show that Axiom 1.(4) implies that, for any  $A \in \Sigma$ , all  $u_A$  are positive affine transformations of  $u_\Omega$ , which means that we can assume  $u_\Omega = u_A$  for all  $A \in \Sigma$ . Define  $u : X \rightarrow \mathbb{R}$  as  $u = u_\Omega$  and  $\pi = \pi_\Omega$ .

**Claim 1.** For any  $A, B \in \Sigma$ ,  $A \supseteq B$ , if  $(\Omega \setminus B)$  is  $\succeq_A$ -null, then  $\succeq_A = \succeq_B$ .

*Proof.* Consider  $A, B \in \Sigma$ ,  $A \supseteq B$  such that  $(\Omega \setminus B)$  is  $\succeq_A$ -null. Notice that since  $A = A \cup B$ , then  $(\Omega \setminus B)$  is  $\succeq_{A \cup B}$ -null. At the same time, since  $(\Omega \setminus A)$  is  $\succeq_{A \cup B}$ -null by Axiom 2. But then, Axiom 4 implies  $\succeq_A = \succeq_B$  as sought.  $\square$

**Claim 2.** For any  $A, B \in M$ , if  $\pi_A(B) = 1 = \pi_B(A)$ , then  $\pi_A = \pi_B$ .

*Proof.* Consider any  $A, B \in M$  such that  $\pi_A(B) = 1 = \pi_B(A)$ . Notice that by construction of  $\pi_A$  and  $\pi_B$  we must have  $\pi_A(A \cap B) = 1 = \pi_B(A \cap B)$ . But then, by Claim 1 we must have  $\pi_A = \pi_{A \cap B} = \pi_B$  as sought.  $\square$

Define now the set  $\mathcal{K} \subseteq \Sigma$  as follows:

$$\mathcal{K} := \{A \in \Sigma : A \text{ is } \succeq\text{-null}\} \cup \{A \in \Sigma : \exists f, g \in \mathcal{F} \text{ s.t. } f \succeq_A g \text{ and } g \succ fAg, \text{ or } f \succ_A g \text{ and } g \succeq fAg\}.$$

These are the events after which either Dynamic Consistency does not apply (null-events), or after which it is violated. Define now  $\epsilon$  as  $\epsilon := \max\{\pi(B) : B \in \mathcal{K}\}$  if  $\mathcal{K} \neq \emptyset$ ,  $\epsilon = 0$  if  $\mathcal{K} = \emptyset$ . Notice also that, by construction of  $\mathcal{K}$ , this implies that  $\epsilon \geq \pi(A)$  for all  $A \in \mathcal{K}$ . Also, by construction we must have  $\epsilon \in [0, 1)$ . To see why, consider any  $A \in \Sigma$  such that  $\pi(A) = 1$ . This implies that  $(\Omega \setminus A)$  is  $\succeq$ -null. But then,  $\succeq_A = \succeq_\Omega$  by Claim 1, hence  $A \notin \mathcal{K}$ . This implies that  $\max\{\pi(B) : B \in \mathcal{K}\} < 1$ , hence  $\epsilon \in [0, 1)$ .

Consider now  $A \in \Sigma \setminus \mathcal{K}$  (notice that this set includes all  $A \in \Sigma$  such that  $\pi(A) > \epsilon$ , by construction of  $\epsilon$ ).

**Claim 3.** For any  $f, g, h \in \mathcal{F}$ ,  $A \in \Sigma \setminus \mathcal{K}$  we have

$$f \succeq_{I_t, A} g \Leftrightarrow fAh \succeq_{I_t} gAh$$

*Proof.* Consider  $f, g, h \in \mathcal{F}$ , and  $A \in \Sigma \setminus \mathcal{K}$ . By construction of  $\mathcal{K}$ , for any  $r, s \in \mathcal{F}$  we have  $rAs \succeq s$  iff  $r \succeq_A s$ . Notice now that by Axiom 2 we have  $fAh \sim_A f$  and  $gAh \sim_A g$ . This implies  $f \succeq_A g$  iff  $fAh \succeq_A gAh$ . Define  $f' := fAh$  and  $g' := gAh$ . Notice that we have  $fAh \succeq_A gAh$  iff  $f' \succeq_A g'$  iff  $f'Ag' \succeq_{I_t} g'$  iff  $fAh \succeq_{I_t} gAh$  (where the last passages use the fact that  $\pi(A) > \epsilon$ ).  $\square$

**Claim 4.** For any  $A \in \Sigma \setminus \mathcal{K}$ ,  $\pi_A(B) = \frac{\pi(A \cup B)}{\pi(A)} = \text{BU}(\pi, A)$ .

*Proof.* Now first of all that for any  $A \in \Sigma \setminus \mathcal{K}$  we must have  $\pi(A) > 0$ , because any  $A$  which is  $\succeq$ -null belongs to  $\mathcal{K}$ . We then have that for any  $f, g \in \mathcal{F}$

$$\begin{aligned} f \succeq_A g &\Leftrightarrow fAh \succeq gAh \\ &\Leftrightarrow \sum_{\omega \in A} \pi(\omega) \mathbb{E}_{f(\omega)}(u) + \sum_{\omega \in \Omega \setminus A} \pi(\omega) \mathbb{E}_{f(\omega)}(u) \geq \sum_{\omega \in A} \pi(\omega) \mathbb{E}_{g(\omega)}(u) + \sum_{\omega \in \Omega \setminus A} \pi(\omega) \mathbb{E}_{g(\omega)}(u) \\ &\Leftrightarrow \sum_{\omega \in A} \pi(\omega) \mathbb{E}_{f(\omega)}(u) \geq \sum_{\omega \in A} \pi(\omega) \mathbb{E}_{g(\omega)}(u) \\ &\Leftrightarrow \frac{1}{\pi(A)} \sum_{\omega \in A} \pi(\omega) \mathbb{E}_{f(\omega)}(u) \geq \frac{1}{\pi(A)} \sum_{\omega \in A} \pi(\omega) \mathbb{E}_{g(\omega)}(u) \end{aligned}$$

Since  $\pi_A$  is unique, this proves that for any  $A \in \Sigma \setminus \mathcal{K}$ ,  $\pi_A(B) = \frac{\pi(A \cup B)}{\pi(A)}$  as sought. In particular, this is also true for any  $A \in \Sigma$  such that  $\pi(A) > \epsilon$ .  $\square$

Define now the set  $\mathcal{K}^*$  as follows.

$$\mathcal{K}^* := \{A \in \Sigma : \pi_\Omega(B) \geq \pi_\Omega(A) \text{ for some } B \in \mathcal{K}\}.$$

Notice that we must have  $\mathcal{K}^* = \{A \in \Sigma : \pi(A) \leq \epsilon\}$  by construction of  $\epsilon$ . This is the set of events such that there exist an event in  $\mathcal{K}$  that is more likely than some of them. Define now the sets  $H := \mathcal{K}^* \cup \{\Omega\}$  and  $M := \{\pi_m \in \Delta(\Omega) : m \in \mathcal{K}^*\} \cup \{\pi_\Omega\}$ .

**Claim 5.** The following holds for no  $A, B, C, D \in \mathcal{K}^* \cup \{\Omega\}$ :  $\pi_A = \pi_B$ ,  $\pi_C = \pi_D$ ,  $\pi_C(A) = 1$ ,  $\pi_A(C) < 1$ ,  $\pi_B(D) = 1$  and  $\pi_D(B) < 1$ .

*Proof.* Say by means of contradiction that such  $A, B, C, D \in M$  and notice that  $\pi_C = \pi_D$  and  $\pi_C(A) = 1$  imply  $\pi_D(A) = 1$ . Similarly,  $\pi_A = \pi_B$  and  $\pi_B(D) = 1$  imply  $\pi_A(D) = 1$ . But then, Claim 2 implies  $\pi_A = \pi_D$ , which means  $\pi_A = \pi_C$ . But by construction of  $\pi_C$  this contradicts  $\pi_A(C) < 1$ .  $\square$

Define now the binary relation  $\triangleright$  on  $M$  as

$$\pi_m \triangleright \pi_{m'} \Leftrightarrow \pi_{m'}(m) = 1 \text{ and } \pi_m(m') < 1.$$

Notice that  $\triangleright$  is well defined by Claim 5, and it is also irreflexive for the same reason.

**Claim 6.** Consider  $\pi_1, \dots, \pi_n \in M$  such that  $\pi_1 \triangleright \dots \triangleright \pi_n$ . Then, there exist  $A_1, \dots, A_n \in \mathcal{K}^*$  such that  $\pi_i = \pi_{A_i}$ ,  $\pi_{A_{i+1}}(A_i) = 1$  and  $\pi_{A_i}(A_{i+1}) < 1$  for  $i = 1, \dots, (n-1)$ .

*Proof.* For simplicity for focus on the case in which  $n = 3$ : it is trivial to show that the proof extends to the general case. Consider  $\pi_1, \pi_2, \pi_3 \in M$  such that  $\pi_1 \triangleright \pi_2 \triangleright \pi_3$ . By construction of  $\triangleright$  we know that there exist  $A_1, A_2, A'_2, A_3 \in \mathcal{K}^*$  such that  $\pi_1 = \pi_{A_1}$ ,  $\pi_2 = \pi_{A_2} = \pi_{A'_2}$ ,  $\pi_3 = \pi_{A_3}$  and  $\pi_{A_2}(A_1) = 1$ ,  $\pi_{A_1}(A_2) < 1$ ,  $\pi_{A_3}(A'_2) = 1$ ,  $\pi_{A'_2}(A_3) < 1$ . Since  $\pi_{A_2} = \pi_{A'_2}$ , we must also have  $\pi_{A'_2}(A_1) = 1$ . If we can prove that we also have  $\pi_{A_1}(A'_2) < 1$ , then we are done. Say, by contradiction, that  $\pi_{A_1}(A'_2) = 1$ . But then, since  $\pi_{A'_2}(A_1) = 1$ , by Claim 2 we have  $\pi_{A_1} = \pi_{A'_2}$ , hence  $\pi_1 = \pi_2$ , which contradicts the fact that  $\triangleright$  is irreflexive.  $\square$

**Claim 7.**  $\triangleright$  is acyclic.

*Proof.* By means of contradiction consider  $\pi_1, \dots, \pi_n \in M$  such that  $\pi_1 \triangleright \dots \triangleright \pi_n \triangleright \pi_1$ . By Claim 6 there exist  $A_1, \dots, A_n \in \mathcal{K}^*$  such that  $\pi_i = \pi_{A_i}$ ,  $\pi_{A_{i+1}}(A_i) = 1$  and  $\pi_{A_i}(A_{i+1}) < 1$  for  $i = 1, \dots, (n-1)$ , and  $\pi_{A_1}(A_n) = 1$  and  $\pi_{A_n}(A_1) < 1$ . Since  $\triangleright$  is irreflexive notice that we must have that  $\pi_{A_i} \neq \pi_{A_{i+1}}$  for  $i = 1, \dots, (n-1)$ , and  $\pi_{A_1} \neq \pi_{A_n}$ . Now construct  $E_i = A_i \cap A_{i+1}$  for  $i = 1, \dots, n$ , and  $E_n = A_n \cap A_1$ . Notice that  $E_i \cup E_{i+1} = (A_{i-1} \cap A_i) \cup (A_i \cap A_{i+1})$ , and since  $\pi_{A_i}(A_{i+1}) = 1$ , then  $\pi_{A_i}(E_i \cup E_{i+1}) = 1$ , hence  $\pi_{A_i} = \pi_{E_i \cup E_{i+1}}$  for  $i = 1, \dots, (n-1)$  by Claim 1. A similar argument shows that  $\pi_{A_n} = \pi_{E_n \cup E_1}$ . Then we have  $\pi_{E_i \cup E_{i+1}}(E_{i+1}) = 1$  for  $i = 1, \dots, (n-1)$ , and  $\pi_{E_n \cup E_1}(E_1) = 1$ . Therefore we have  $\pi_{A_i} = \pi_{E_i \cup E_{i+1}} = \pi_{E_{i+1}}$  for  $i = 1, \dots, (n-1)$ , and  $\pi_{E_n \cup E_1} = \pi_{E_1}$ . Hence  $\pi_{A_i} = \pi_{E_{i+1}}$  for  $i = 1, \dots, (n-1)$ , and  $\pi_{A_n} = \pi_{E_1}$ . At the same time, since  $\pi_{A_i} \neq \pi_{A_{i+1}}$  for  $i = 1, \dots, (n-1)$  and  $\pi_{A_1} \neq \pi_{A_n}$ , then  $\pi_{E_i} \neq \pi_{E_{i+1}}$  for  $i = 1, \dots, (n-1)$ , and  $\pi_{E_1} \neq \pi_{E_n}$ . But this contradicts Axiom 4.  $\square$

Define  $\gamma_* := \max\{\pi_A(B) : A, B \in \mathcal{K}^*, \pi_A(B) < 1\}$  and define  $\delta := \frac{1}{|M|} \frac{1-\gamma_*}{\gamma_*}$ . (Notice that  $\gamma_*$  is well defined since  $M$  is finite, and that  $\gamma_* \in (0, 1)$  and consequently  $\delta > 0$ .) Notice also that since  $M$  is finite and  $\triangleright$  is acyclic. Consider now the transitive closure of  $\triangleright$  and call it  $\hat{\triangleright}$ . Since  $\triangleright$  is irreflexive and antisymmetric, so must  $\hat{\triangleright}$ . To see why, say, by contradiction, that we have  $A, B \in M$  such that  $A \hat{\triangleright} B \hat{\triangleright} A$ . Then, by definition of transitive closure, there must exist  $m_1, \dots, m_{n+m} \in M$  such that  $A \triangleright m_1 \triangleright \dots \triangleright m_n \triangleright B$  and  $B \triangleright m_{n+1} \triangleright \dots \triangleright m_{n+m} \triangleright A$ . But this violates the acyclicity of  $\triangleright$ .

Since  $M$  is finite, enumerate it and construct the function  $f$  as follows. Set  $f(m_1) = 0$ . Consider  $m_n$ . Assign to  $f(m_n)$  any value such that: for all  $i < n$ ,  $f(m_n) \neq f(m_i)$ ;  $f(m_n) > f(m_i)$  if  $m_n \hat{\triangleright} m_i$ ; and  $f(m_n) < f(m_i)$  if  $m_i \hat{\triangleright} m_n$ . To see why this is always possible, notice that for all  $m_n, m_i, m_j$ , with  $n \geq i, j$ , if we have  $m_n \hat{\triangleright} m_i$  and  $m_j \hat{\triangleright} m_n$ , then we must also have that  $m_j \hat{\triangleright} m_i$  since  $\hat{\triangleright}$  is transitive, which implies that we must also have  $f(m_j) > f(m_i)$ . Thus  $f$  is well defined. Normalize now the function  $f$  so that it has a range in  $(0, \delta)$  and call it  $v$ . Notice that we must have that  $\pi \triangleright \pi'$  implies  $v(\pi) > v(\pi')$ , and such that  $v(\pi) \neq v(\pi')$  for all  $\pi, \pi' \in M$ .

Construct now  $\rho \in \Delta(\Delta(\Omega))$  as

$$\rho(\pi) := \frac{v(\pi) + \frac{1}{|M|}}{\sum_{m \in M} v(\pi) + \frac{1}{|M|}}.$$

for all  $\pi \in M$ ,  $\rho(\pi) = 0$  otherwise.

We now turn to show that  $M$  and  $\rho$  that we just constructed are the ones that we are looking for.

**Claim 8.**  $\{\pi_\Omega\} = \arg \max_{\pi \in \Delta(\Omega)} \rho(\pi)$ .

*Proof.* Say, by means of contradiction that there exist  $\pi \in \Delta(\Omega)$  such that  $\pi \neq \pi_\Omega$  and  $\rho(\pi) \geq \rho(\pi_\Omega)$ . For this to be possible we must have  $\rho(\pi) > 0$ , which in turns implies that  $\pi = \pi_m$  for some  $m \in H$ . Say first that we have that  $\pi_\Omega(m) = 1$  for all  $m \in H$  such that  $\pi_m = \pi$ . Then, since we also have that  $\pi_m(\Omega) = 1$  for all  $m \in H$  such that  $\pi_m = \pi$ , then Claim 2 would imply  $\pi = \pi_\Omega$ , a contradiction. This means that we must have  $\pi_\Omega(m) < 1$  for some  $m \in H$  such that  $\pi_m = \pi$ . Also, notice that we must have  $\pi_m(\Omega) = 1$  (by definition),

which implies that we have  $\pi_\Omega \triangleright \pi_m = \pi$ . But then we must have  $v(\pi_\Omega) > v(\pi)$  by construction of  $v$ , which in turns implies that we must have  $\rho(\pi_\Omega) > \rho(\pi)$ , a contradiction.  $\square$

Notice now that for any  $A \in \Sigma$ , given the definition of  $\epsilon$ , the agent will behave as prescribed in the representation if  $\pi(A) > \epsilon$ . We now turn to analyze the events with probability below the threshold.

**Claim 9.** For any  $A \in H$ , if  $\pi \in \arg \max_{m \in M} \text{BU}(\rho, A)$ , then  $\pi(A) = 1$ .

*Proof.* Consider  $A \in H$  and say by means of contradiction that there exist  $\pi \in \arg \max_{m \in M} \text{BU}(\rho|A)$  such that  $\pi(A) < 1$ . This means that we must have  $\pi(A)\rho(\pi) \geq \pi_A(A)\rho(\pi_A)$ . Notice that since  $A \in H$ , by construction we must have  $\pi_A(A) = 1$ , which implies that we must have  $\pi(A)\rho(\pi) \geq \rho(\pi_A)$ . By construction of  $\rho$  this is possible only if  $\pi \in M$  and

$$(v(\pi) + \frac{1}{|M|})\pi_A \geq v(\pi_A) + \frac{1}{|M|}.$$

Since by construction  $v(\pi_A) > 0$  and  $\pi_A \leq \gamma_*$ , then this implies

$$(v(\pi) + \frac{1}{|M|})\gamma_* > \frac{1}{|M|} \Rightarrow v(\pi) > \frac{1}{|M|} \frac{1 - \gamma_*}{\gamma_*}.$$

But since  $v$  has range  $(0, \frac{1}{|M|} \frac{1 - \gamma_*}{\gamma_*})$ , this is a contradiction.  $\square$

**Claim 10.** For any  $A \in H$  we have  $\{\pi_A\} = \arg \max_{m \in M} \text{BU}(\rho, A)$ .

*Proof.* Consider  $A \in H$  and say, by means of contradiction, that we have  $\pi \in \arg \max_{m \in M} \text{BU}(\rho|A)$  for some  $\pi \neq \pi_A$ . This means that we have  $\pi(A)\rho(\pi) \geq \pi_A(A)\rho(\pi_A)$ . By Claim 9 we know that we must have  $\pi(A) = 1$  and since  $\pi_A(A) = 1$ , then this means that we must have  $\rho(\pi) \geq \rho(\pi_A)$ . For this to be possible we must have  $\pi \in M$ , which implies that there exists  $B \in H \cup \{\Omega\}$  such that  $\pi_B = \pi$ . Now, if we have  $\pi_A(B) = 1$ , then since  $\pi_B(A) = \pi(A) = 1$ , by Claim 2 we have  $\pi = \pi_A$ , a contradiction. Therefore, we must have  $\pi_A(B) < 1$ . But then, we must have  $\pi_A \triangleright \pi_B = \pi$ , which implies  $v(\pi_A) > v(\pi_B)$ , hence  $\rho(\pi_A) > \rho(\pi)$ , a contradiction.  $\square$

Claim 10 proves the representation. Notice also that point (4) of the representation is trivially true by our construction of  $\rho$ . We have therefore found an Hypothesis Testing Representation  $(u, \rho, \epsilon)$ . Notice that by the definition of  $\epsilon$ ,  $u, \rho, \epsilon$  is also a minimal representation.

[*Necessity of the Axioms*] The proof the necessity of Axiom 1 (WbP) and Axiom is standard practice. Axiom 2 is immediate from the representation. We are left with Axiom 4. Consider  $A_1, \dots, A_n \in \Sigma$  such that  $\succeq_{A_1} \neq \succeq_{A_n}$ ,  $(\Omega \setminus A_{i+1})$  is  $\succeq_{(A_i \cup A_{i+1})}$ -null for  $i = 1, \dots, (n-1)$ , and  $(\Omega \setminus A_1)$  is  $\succeq_{(A_n \cup A_1)}$ -null. Consider first the case in which  $\pi(A_i \cup A_{i+1}) > \epsilon$  for  $i = 1, \dots, (n-1)$ , and  $\pi(A_1 \cup A_n) > \epsilon$ . By the representation, this means that we must have  $\pi_{A_i \cup A_{i+1}} = \text{BU}(\pi, A_i \cup A_{i+1})$  for  $i = 1, \dots, (n-1)$ , and  $\pi_{A_1 \cup A_n} = \text{BU}(\pi, A_1 \cup A_n)$ . But since  $(\Omega \setminus A_{i+1})$  is  $\succeq_{(A_i \cup A_{i+1})}$ -null for  $i = 1, \dots, (n-1)$ , hence  $\pi_{A_i \cup A_{i+1}}(A_i \setminus A_{i+1}) = 0$ , then by construction of BU we must have  $\pi(A_i \setminus A_{i+1}) = 0$  for  $i = 1, \dots, n$ . For the same reason we must also have  $\pi(A_n \setminus A_1) = 0$ . But then we must have that  $\pi(\bigcap_{i=1}^n A_i) = \pi(\bigcup_{i=1}^n A_i)$ , and so, by definition of BU, we must have  $\pi_{A_1} = \pi_{A_n}$ , hence  $\succeq_{A_1} = \succeq_{A_n}$ , a contradiction.

Consider now the more general case in which there exist some  $i$  such that  $\pi(A_i \cup A_{i+1}) > \epsilon$ , or in which we have  $\pi(A_1 \cup A_n) > \epsilon$ . Say without loss of generality that we have  $\pi(A_1 \cup A_2) > \epsilon$ . By the representation it must be  $\pi(A_1 \cup A_2) = \text{BU}(\pi, A_1 \cup A_2)$ . At the same time we have  $\pi(A_1 \cup A_2)(A_1 \setminus A_2) = 0$ , and by definition of BU this means that we have  $\pi(A_1 \cup A_2) = \pi(A_2)$ . Then  $\pi(A_2) > \epsilon$ , hence  $\pi(A_2 \cup A_3) > \epsilon$ . Proceed like this to prove that we must have  $\pi(A_i \cup A_{i+1}) > \epsilon$  for  $i = 1, \dots, (n-1)$ , and that  $\pi(A_1 \cup A_n) > \epsilon$ . But we have already shown that this leads to a contradiction.

We are left with the case in which  $\pi(A_i \cup A_{i+1}) \leq \epsilon$  for  $i = 1, \dots, (n-1)$  and  $\pi(A_1 \cup A_n) \leq \epsilon$ . Now define  $E_1, \dots, E_n$  as  $E_i = A_i \cup A_{i+1}$  for  $i = 1, \dots, (n-1)$  and  $E_n = A_1 \cup A_n$ . Notice that we must have  $\pi(E_i) \leq \epsilon$  for  $i = 1, \dots, n$ . Moreover, by construction we must also have that  $\text{BU}(\pi_{E_i}^*, E_i)(E_{i+1}) = 1$  for  $i = 1, \dots, n-1$  and  $\text{BU}(\pi_{E_n}^*, E_n)(E_1) = 1$ . This implies  $\pi_{E_i}^*(E_{i+1}) \geq \pi_{E_i}^*(E_i)$  for  $i = 1, \dots, n-1$  and  $\pi_{E_n}^*(E_1) \geq \pi_{E_n}^*(E_n)$ . Now, notice that since  $\pi_{E_i}^*$  is the unique element in  $\arg \max_{\pi \in \Delta(\Omega)} \text{BU}(\rho, E_i)(\pi)$  for  $i = 1, \dots, n$ , then we must have that for all  $i, j = 1, \dots, n$ ,  $\rho(\pi_{E_i}^*)\pi_{E_i}^*(E_i) > \rho(\pi_{E_j}^*)\pi_{E_j}^*(E_i)$  if  $\pi_{E_i}^* \neq \pi_{E_j}^*$ . This, together with the fact that  $\pi_{E_i}^*(E_{i+1}) \geq \pi_{E_i}^*(E_i)$  for  $i = 1, \dots, n-1$  and  $\pi_{E_n}^*(E_1) \geq \pi_{E_n}^*(E_n)$ , implies that  $\rho(\pi_{E_i}^*)\pi_{E_{i+1}}^*(E_{i+1}) \geq \rho(\pi_{E_{i+1}}^*)\pi_{E_{i+1}}^*(E_{i+1})$ , where the inequality is strict if  $\pi_{E_i}^* \neq \pi_{E_{i+1}}^*$ . Hence for  $i = 1, \dots, (n-1)$ ,  $\rho(\pi_{E_i}^*) \geq \rho(\pi_{E_{i+1}}^*)$ , where the inequality is strict if  $\pi_{E_i}^* \neq \pi_{E_{i+1}}^*$ , and  $\rho(\pi_{E_n}^*) \geq \rho(\pi_{E_1}^*)$ , where the inequality is strict if  $\pi_{E_n}^* \neq \pi_{E_1}^*$ . But then we have  $\rho(\pi_{E_1}^*) \geq \rho(\pi_{E_2}^*) \geq \dots \geq \rho(\pi_{E_n}^*) \geq \rho(\pi_{E_1}^*)$ , and so none of these inequalities can be strict, hence  $\pi_{E_1}^* = \pi_{E_i}^*$  for  $i = 1, \dots, n$ . But then,  $\text{BU}(\pi_{E_1}^*, E_i)(E_{i+1}) = 1$  for  $i = 1, \dots, n-1$  and  $\text{BU}(\pi_{E_1}^*, E_n)(E_1) = 1$ , which implies that  $\text{BU}(\pi_{E_1}^*, E_n) = \text{BU}(\pi_{E_1}^*, E_1)$ , hence  $\pi_{E_1} = \pi_{E_n}$  and  $\succeq_{E_1} = \succeq_{E_n}$ . Repeating the same argument we obtain  $\succeq_{E_i} = \succeq_{E_{i+1}}$  for  $i = 1, \dots, (n-1)$ , hence  $\succeq_{E_i} = \succeq_{E_j}$  for  $i, j = 1, \dots, n$ . Now recall since  $(\Omega \setminus A_n)$  is  $\succeq_{(A_{n-1} \cup A_n)}$ -null, and notice now that for any  $\pi \in \arg \max_{\pi \in \Delta(\Omega)} \text{BU}(\rho, A_{n-1} \cup A_n)$  such that  $\text{BU}(\pi, A_{n-1} \cup A_n)(A_n) = 1$  (since  $(\Omega \setminus A_n)$  is  $\succeq_{(A_{n-1} \cup A_n)}$ -null), we must also have  $\pi \in \arg \max_{\pi \in \Delta(\Omega)} \text{BU}(\rho, A_n)$  and  $\text{BU}(\pi, A_n) = \text{BU}(\pi, A_{n-1} \cup A_n)$  (by definition of BU). But since the argmax is unique, then  $\pi_{A_n} = \pi_{A_{n-1} \cup A_n} = \pi_{E_{n-1}}$ . Similarly we obtain that  $\pi_{A_1} = \pi_{A_1 \cup A_n} = \pi_{E_n}$ . But since  $\pi_{E_{n-1}} = \pi_{E_n}$ , then  $\pi_{A_1} = \pi_{A_n}$ , hence  $\succeq_{A_1} = \succeq_{A_n}$ , a contradiction.

[ $\epsilon = 0$  iff *Dynamic Consistency*] Notice first of all that if  $\epsilon = 0$ , then the agent updates her prior using Bayes' rule every time she is told that a non-null event has occurred, which it is well known to imply that Dynamic Consistency is satisfied. Consider now the case in which the  $\{\succeq_A\}_{A \in \Sigma}$  satisfies Dynamic Consistency. Let us say, by means of contradiction, that we have a minimal Hypothesis Testing representation  $(u, \rho, \epsilon)$  of  $\{\succeq_A\}_{A \in \Sigma}$  in which  $\epsilon \neq 0$ . Since  $\{\succeq_A\}_{A \in \Sigma}$  satisfies Dynamic Consistency, however, then  $(u, \rho, 0)$  must also represent it, contradicting the minimality of  $(u, \rho, \epsilon)$ . Q.E.D.

### Proof of Theorem 2

[*Sufficiency of the Axioms*] We proceed in a similar way to how we proceeded for the proof of Theorem 1. Given Axioms 5, from Gilboa and Schmeidler (1989) we know that for any  $A \in \Sigma$ , there exist  $u_A : X \rightarrow \mathbb{R}$ ,  $\Pi_A \subseteq \Delta(\Omega)$ ,  $\Pi$  convex and compact, such that for any  $f, g \in \mathcal{F}$

$$f \succeq_A g \Leftrightarrow \min_{\pi \in \Pi_A} \sum_{\omega \in A} \pi(\omega) \mathbb{E}_{f(\omega)}(u) \geq \min_{\pi \in \Pi_A} \sum_{\omega \in A} \pi(\omega) \mathbb{E}_{g(\omega)}(u) \quad (\text{A.2})$$

where  $\Pi_A$  is unique and  $u_A$  is unique up to a positive affine transformation. It is also standard practice to show that Axiom 5(5) implies that, for any  $A \in \Sigma$ , all  $u_A$  are positive affine transformations of  $u_\Omega$ , which means that we can assume  $u_\Omega = u_A$  for all  $A \in \Sigma$ . Define  $u : X \rightarrow \mathbb{R}$  as  $u = u_\Omega$  and  $\pi = \pi_\Omega$ . Moreover, notice that for any  $A, B \in \Sigma$   $A$  is  $\succeq_B$ -null if and only if  $\pi(A) = 0$  for all  $\pi \in \Pi_B$ .

Notice that Claims 1 holds true here as well. Moreover, notice the following claim (which parallels Claim 2 in the proof of Theorem 1.)

**Claim 11.** For any  $A, B \in M$ , if  $\pi_A(B) = 1 = \pi_B(A)$  for all  $\pi_A \in \Pi_A$  and  $\pi_B \in \Pi_B$ , then  $\Pi_A = \Pi_B$ .

*Proof.* Consider any  $A, B \in M$  such that  $\pi_A(B) = 1 = \pi_B(A)$  for all  $\pi_A \in \Pi_A$  and  $\pi_B \in \Pi_B$ . Notice that by construction of  $\Pi_A$  and  $\Pi_B$  we must have  $\pi_A(A \cap B) = 1 = \pi_B(A \cap B)$  for all  $\pi_A \in \Pi_A$  and  $\pi_B \in \Pi_B$ . Hence  $(\Omega \setminus (A \cap B))$  is both  $\succeq_A$ -null and  $\succeq_B$ -null. But then, by Claim 1 we must have  $\Pi_A = \Pi_{A \cap B} = \Pi_B$  as sought. □

Define now the set  $\mathcal{K}_{AA} \subseteq \Sigma$  as follows:

$\mathcal{K}_{AA} := \{A \in \Sigma : A \text{ is } \succeq\text{-null}\} \cup \{A \in \Sigma : \exists f, g \in \mathcal{F} \text{ s.t. } f \succeq_A^* g \text{ and } g \succ^* fAg, \text{ or } f \succ_A^* g \text{ and } g \succeq^* fAg\}$ .

These are the events after which either Reduced Dynamic Consistency (Axiom 6) does not apply (null-events), or after which it is violated. Define  $\epsilon$  as  $\epsilon := \max_{A \in \mathcal{K}} \max_{\pi \in \Pi_\Omega} \pi(A)$  if  $\mathcal{K}_{AA} \neq \emptyset$ ,  $\epsilon = 0$  if  $\mathcal{K}_{AA}^* = \emptyset$ . (This is well defined since  $\Omega$  is finite.)

Consider now  $A \in \Sigma \setminus \mathcal{K}_{AA}$  (notice that this set includes all  $A \in \Sigma$  such that  $\pi(A) > \epsilon$  for all  $\pi \in \Pi_\Omega$ , by construction of  $\epsilon$ ).

**Claim 12.** For any  $A \in \Sigma \setminus \mathcal{K}_{AA}$ ,  $f, g, h \in \mathcal{F}$ , we have

$$f \succeq_A g \Leftrightarrow fAh \succeq_\Omega gAh$$

*Proof.* Consider  $f, g, h \in \mathcal{F}$ ,  $A \notin \mathcal{K}_{AA}^*$ . Notice first of all that by Axiom 2 we have  $fAh \sim_A f$  and  $gAh \sim_A g$ . This implies that we have  $f \succeq_A g$  iff  $fAh \succeq_A gAh$ . Define  $f' := fAh$  and  $g' := gAh$ . Notice that we have  $fAh \succeq_A gAh$  iff  $f' \succeq_A g'$  iff (since  $A \notin \mathcal{K}_{AA}^*$ )  $f'Ag' \succeq_\Omega g'$  iff  $fAh \succeq_\Omega gAh$ , as sought.  $\square$

Finally, from Ghirardato, Maccheroni, and Marinacci (2004) we know that for any  $A \in \Sigma$ ,  $\succeq_A^*$  satisfies monotonicity, continuity and independence, and it can be represented by

$$f \succeq_A^* g \Leftrightarrow \sum_{\omega \in \Omega} \pi(\omega) \mathbb{E}_{f(\omega)}(u) \geq \sum_{\omega \in \Omega} \pi(\omega) \mathbb{E}_{g(\omega)}(u) \quad \forall \pi \in \Pi_A$$

where  $\Pi_A$  is a compact and convex subset of  $\Delta(\Omega)$ , it is the same as the one in Equation A.2, and it is unique.<sup>18</sup> This means that for any  $A \in \Sigma$  such that  $\pi(A) > \epsilon$  for some  $\pi \in \Pi_\Omega$  (which also means  $\pi(A) > 0$  for some  $\pi \in \Pi_\Omega$ ) we have

$$\begin{aligned} f \succeq_A^* g &\Leftrightarrow fAh \succeq_\Omega^* gAh \\ &\Leftrightarrow \sum_{\omega \in A} \pi(\omega) \mathbb{E}_{f(\omega)}(u) + \sum_{\omega \in \Omega \setminus A} \pi(\omega) \mathbb{E}_{f(\omega)}(u) \geq \sum_{\omega \in A} \pi(\omega) \mathbb{E}_{g(\omega)}(u) + \sum_{\omega \in \Omega \setminus A} \pi(\omega) \mathbb{E}_{g(\omega)}(u) \quad \forall \pi \in \Pi_\Omega \\ &\Leftrightarrow \sum_{\omega \in A} \pi(\omega) \mathbb{E}_{f(\omega)}(u) \geq \sum_{\omega \in A} \pi(\omega) \mathbb{E}_{g(\omega)}(u) \quad \forall \pi \in \Pi_\Omega \\ &\Leftrightarrow \frac{1}{\pi(A)} \sum_{\omega \in A} \pi(\omega) \mathbb{E}_{f(\omega)}(u) \geq \frac{1}{\pi(A)} \sum_{\omega \in A} \pi(\omega) \mathbb{E}_{g(\omega)}(u) \quad \forall \pi \in \Pi_\Omega \end{aligned}$$

Since  $\Pi_A$  is unique, this proves that for any  $A \in \Sigma$  such that  $\pi(A) > \epsilon$  for some  $\pi \in \Pi_\Omega$ , then  $\Pi_A(B) = \hat{B}U(\Pi_\Omega, A)$ .

Define now the set  $\mathcal{K}^*$  as follows.

$$\mathcal{K}_{AA}^* := \{A \in \Sigma : \pi(A) \leq \epsilon \text{ for some } \pi \in \Pi\}.$$

Define the sets  $H_{AA} := \mathcal{K}_{AA}^* \cup \{\Omega\}$  and  $M_{AA} := \{\pi_m \in \Delta(\Omega) : m \in \mathcal{K}_{AA}^*\} \cup \{\pi_\Omega\}$ .

We can now proceed replicating exactly the steps in the proof of Theorem 1 and prove the adaptations of Claims 5, 6, and 7, with the following modifications: use sets of priors  $\Pi_m$  instead of priors  $\pi_m$ ; whenever

<sup>18</sup>See Ghirardato, Maccheroni, and Marinacci (2004). In Section 5.1 they discuss how their Theorem 14 implies that the set of priors found by the representation of  $\succeq^*$  using their Theorem 11 must coincide with the one found with a representation of  $\succeq$  a la Gilboa and Schmeidler (1989).

we have  $\pi_m(A) = 1$  replace it with  $\pi(A) = 1$  for all  $\pi \in \Pi_m$ ; replace the conditions  $\pi(A) \leq \epsilon$  with the corresponding condition  $\pi(A) \leq \epsilon$  for all  $\pi \in \Pi_\Omega$ ; use the set  $\mathcal{K}_{AA}^*, H_{AA}, M_{AA}$  instead of  $\mathcal{K}^*, H, M$ . In particular, construct the preference  $\triangleright$  on  $M_{AA}$ .

Define  $\gamma_{AA}^* := \max\{\min_{\pi \in \Pi_A} \pi(B) : A, B \in \mathcal{K}_{AA}^*, \pi(B) < 1 \text{ for some } \pi \in \Pi_A\}$  and define  $\delta := \frac{1}{|M|} \frac{1-\gamma_*}{\gamma_*}$ . (Notice that  $\gamma_*$  is well defined since  $M$  is finite, and that  $\gamma_* \in (0, 1)$  and consequently  $\delta > 0$ .) Proceed like in the proof of Theorem 1 in constructing the transitive closure of  $\hat{\triangleright}$  of  $\triangleright$ , and the function  $f$  and  $v$  on  $M_{AA}$ , and construct  $\rho \in \Delta(\mathcal{G})$  as

$$\rho(\Pi) := \frac{v(\Pi) + \frac{1}{|M|}}{\sum_{m \in M} v(\Pi) + \frac{1}{|M|}}.$$

for all  $\Pi \in M$ ,  $\rho(\Pi) := 0$  otherwise. Finally, we need to prove that  $M_{AA}$  and  $\rho$  that we just constructed are the ones that we are looking for. But it's easy to replicate the passages in the proof of Claims 8, 9, 10 in the proof of Theorem 1 (once again whenever we have  $\pi_m(A) = 1$  we need to replace it with  $\pi(A) = 1$  for all  $\pi \in \Pi_m$ , and use  $M_{AA}$ ,  $v$ , and  $\rho$  as constructed here). Condition (4) will also be trivially true here as well by construction, since for any event  $A \in \mathcal{K}_{AA}^*$  we construct  $\Pi \in \text{supp}(\rho)$  such that each  $\pi(A) > 0$  for all  $\pi \in \Pi$ .

[*Necessity of the Axioms*] Axiom 1 and Axiom 2 are immediate. We are left with Axiom 4. Consider  $A_1, \dots, A_n \in \Sigma$  such that  $\succeq_{A_1} \neq \succeq_{A_n}$ ,  $(\Omega \setminus A_{i+1})$  is  $\succeq_{(A_i \cup A_{i+1})}$ -null for  $i = 1, \dots, (n-1)$ , and  $(\Omega \setminus A_1)$  is  $\succeq_{(A_n \cup A_1)}$ -null. Consider first the case in which  $\pi(A_i \cup A_{i+1}) > \epsilon$  for some  $\pi \in \Pi$ , for all  $i = 1, \dots, (n-1)$ , and  $\pi(A_1 \cup A_n) > \epsilon$  for some  $\pi \in \Pi$ . Since  $(\Omega \setminus A_{i+1})$  is  $\succeq_{(A_i \cup A_{i+1})}$ -null for  $i = 1, \dots, (n-1)$ , then it must be that  $\pi(A_i \setminus A_{i+1}) = 0$  for all  $\pi \in \Pi_{A_i \cup A_{i+1}}$  for  $i = 1, \dots, (n-1)$ , and  $\pi(A_n \setminus A_1) = 0$  for all  $\pi \in \Pi_{A_n \cup A_1}$ . Notice that by the representation it must be that  $\Pi_{A_i \cup A_{i+1}} = \text{BU}(\Pi_\Omega, A_i \cup A_{i+1})$  for  $i = 1, \dots, (n-1)$ , and  $\Pi_{A_1 \cup A_n} = \text{BU}(\Pi_\Omega, A_1 \cup A_n)$ . But then, for all  $\pi \in \Pi_\Omega$  we have  $\pi(A_i \setminus A_{i+1}) = 0$  for  $i = 1, \dots, (n-1)$ , and  $\pi(A_1 \setminus A_n) = 0$ . Hence  $\pi(\cup_{i=1}^n A_i) = \pi(\cap_{i=1}^n A_i)$  for all  $\pi \in \Pi_\Omega$ , and so  $\Pi = \text{BU}(\Pi_\Omega, \cap_{i=1}^n A_i) = \text{BU}(\Pi_\Omega, A_j \cup A_{j+1}) = \text{BU}(\Pi_\Omega, A_1 \cup A_n)$  for  $j = 1, \dots, n-1$ . But this implies  $\succeq_{A_1} \neq \succeq_{A_n}$ , a contradiction.

Consider now the more general case in which there exist some  $i$  such that  $\pi(A_i \cup A_{i+1}) > \epsilon$  for some  $\pi \in \Pi_\Omega$ , or the case in which  $\pi(A_1 \cup A_n) > \epsilon$  for some  $\pi \in \Pi_\Omega$ . Say without loss of generality that we have  $\pi(A_1 \cup A_2) > \epsilon$  for some  $\pi \in \Pi_\Omega$ . By the representation it must be  $\Pi_{A_1 \cup A_2} = \text{BU}(\Pi_\Omega, A_1 \cup A_2)$ . At the same time we have  $\pi(A_1 \cup A_2)(A_1 \setminus A_2) = 0$  for all  $\pi \in \Pi_{A_1 \cup A_2}$ , and by definition of BU this means that we have  $\pi(A_1 \cup A_2) = \pi(A_2)$  for all  $\pi \in \Pi_\Omega$ . But then  $\pi(A_2) > \epsilon$  for some  $\pi \in \Pi_\Omega$ , hence  $\pi(A_2 \cup A_3) > \epsilon$  for some  $\pi \in \Pi_\Omega$ . Proceed like this to prove that we must have  $\pi(A_i \cup A_{i+1}) > \epsilon$  for some  $\pi \in \Pi_\Omega$  for  $i = 1, \dots, (n-1)$ , and that  $\pi(A_1 \cup A_n) > \epsilon$  for some  $\pi \in \Pi_\Omega$ . But we have already shown that this leads to a contradiction.

Then we are left with the case in which  $\pi(A_i \cup A_{i+1}) \leq \epsilon$  for all  $\pi \in \Pi$  for  $i = 1, \dots, (n-1)$ , and  $\pi(A_1 \cup A_n) \leq \epsilon$  for all  $\pi \in \Pi$ . Then, we can replicate the same passages in the proof of the necessity of the axioms in Theorem 1, but putting  $\pi(A) = 1$  for all  $\pi \in \Pi_m$  whenever we have  $\pi_m(A) = 1$  for  $m \in \Sigma$ .

[ $\epsilon = 0$  iff *Reduced Dynamic Consistency*] Ghirardato, Maccheroni, and Marinacci (2008) show that the agent updates her set of priors using Bayes' rule every time she is told that a non-null event has occurred if and only if she satisfies Reduced Dynamic Consistency. If  $\epsilon = 0$ , therefore, Reduced Dynamic Consistency applies. Conversely, assume that Reduced Dynamic Consistency is satisfied and let us say, by means of contradiction, that we have a minimal Hypothesis Testing representation with Ambiguity Aversion  $(u, \rho, \epsilon)$  of  $\{\succeq_A\}_{A \in \Sigma}$  in which  $\epsilon \neq 0$ . Since  $\{\succeq_A\}_{A \in \Sigma}$  satisfies Reduced Dynamic Consistency, however, then  $(u, \rho, 0)$  must also represent it, contradicting the minimality of  $(u, \rho, \epsilon)$ . *Q.E.D.*

## References

- BLUME, L., A. BRANDENBURGER, AND E. DEKEL (1991): "Lexicographic probabilities and choice under uncertainty," *Econometrica: Journal of the Econometric Society*, 61–79.
- BRANDT, M., Q. ZENG, AND L. ZHANG (2004): "Equilibrium stock return dynamics under alternative rules of learning about hidden states," *Journal of Economic Dynamics and Control*, 28, 1925–1954.

- BRAV, A. AND J. HEATON (2002): “Competing theories of financial anomalies,” *Review of Financial Studies*, 15, 575.
- CAMERER, C. (1995): “Individual decision making,” in *Handbook of Experimental Economics*, ed. by A. R. J. Kagel, Princeton University Press, Princeton, vol. 2.
- DE BONDT, W. AND R. THALER (1985): “Does the stock market overreact?” *The Journal of Finance*, 40, 793–805.
- EPSTEIN, L. (2006): “An axiomatic model of non-Bayesian updating,” *Review of Economic Studies*, 73, 413–436.
- EPSTEIN, L. AND M. LE BRETON (1993): “Dynamically Consistent Beliefs Must Be Bayesian,” *Journal of Economic Theory*, 61, 1–22.
- EPSTEIN, L., J. NOOR, AND A. SANDRONI (2008): “Non-Bayesian updating: a theoretical framework,” *Theoretical Economics*, 3, 193–229.
- (2009): “Non-Bayesian Learning,” .
- EPSTEIN, L. AND M. SCHNEIDER (2003): “Recursive multiple-priors,” *Journal of Economic Theory*, 113, 1–31.
- FOSTER, D. AND H. YOUNG (2003): “Learning, hypothesis testing, and Nash equilibrium,” *Games and Economic Behavior*, 45, 73–96.
- GENNAIOLI, N. AND A. SHLEIFER (2009): “What comes to mind,” .
- GHIRARDATO, P. (2002): “Revisiting Savage in a conditional world,” *Economic Theory*, 20, 83–92.
- GHIRARDATO, P., F. MACCHERONI, AND M. MARINACCI (2004): “Differentiating ambiguity and ambiguity attitude,” *Journal of Economic Theory*, 118, 133–173.
- (2008): “Revealed Ambiguity and Its Consequences: Updating,” in *Advances in Decision Making under Risk and Uncertainty. Selected Papers from the FUR 2006 conference*, ed. by M. Abdellaoui and J. Hey, Berlin: Springer-Verlag.
- GILBOA, I. AND L. SAMUELSON (2008): “Preferring Simplicity,” Mimeo Yale University.
- GILBOA, I. AND D. SCHMEIDLER (1989): “Maxmin expected utility with a non-unique prior,” *Journal of Mathematical Economics*, 18, 141–153.
- (1993): “Updating Ambiguous Beliefs,” *Journal of Economic Theory*, 59, 34–49.
- GRETHER, D. (1992): “Testing Bayes rule and the representativeness heuristic: Some experimental evidence,” *Journal of Economic Behavior & Organization*, 17, 31–57.
- GRIFFIN, D. AND A. TVERSKY (1992): “The weighing of evidence and the determinants of confidence,” *Cognitive Psychology*, 24, 411–435.
- HANANY, E. AND P. KLIBANOFF (2006): “Updating Preferences with Multiple Priors,” M, Northwestern University.
- HOLT, C. AND A. SMITH (2009): “An update on Bayesian updating,” *Journal of Economic Behavior and Organization*, 69, 125–134.
- KAHNEMAN, D. AND A. TVERSKY (1982): *Intuitive Prediction: Biases and Corrective Procedures*, London: Cambridge University Press.
- KOCHOV, A. (2009): “A Model of Limited Foresight,” Mimeo, University of Rochester.
- MACCHERONI, F., M. MARINACCI, AND A. RUSTICHINI (2006): “Dynamic variational preferences,” *Journal of Economic Theory*, 128, 4–44.
- MULLAINATHAN, S. (2002): “Thinking through categories,” *NBER working paper*.
- MULLAINATHAN, S., J. SCHWARTZSTEIN, AND A. SHLEIFER (2008): “Coarse Thinking and Persuasion\*,” *Quarterly Journal of Economics*, 123, 577–619.
- MYERSON, R. (1986a): “Axiomatic foundations of bayesian decision theory,” *Center for Mathematical Studies in Economics and Management Science, Northwestern University Working Paper*, 671.
- (1986b): “Multistage games with communication,” *Econometrica*, 54, 323–358.
- RABIN, M. (1998): “Psychology and economics,” *Journal of Economic Literature*, 11–46.
- (2002): “Inference By Believers in the Law of Small Numbers\*,” *Quarterly Journal of Economics*, 117, 775–816.
- SINISCALCHI, M. (2006): “Dynamic Choice under Ambiguity,” Mimeo, Northwestern University.
- TVERSKY, A. AND D. KAHNEMAN (1974): “Judgment under uncertainty: Heuristics and biases,” *Science*, 185, 1124–1131.
- (1981): “The framing of decisions and the psychology of choice,” *Science*, 211, 453–458.