Department of Economics
University of Pittsburgh
Kenneth P. Dietrich School of Arts and Sciences

WORKING PAPER SERIES

16/007

Who is a Bayesian?

Ehud Lehrer and Roee Teper

February 25, 2016

University of Pittsburgh • Dietrich School of Arts and Sciences • Department of Economics
4901 Wesley W. Posvar Hall • 230 South Bouquet Street • Pittsburgh, PA 15260

WHO IS A BAYESIAN?

EHUD LEHRER† AND ROEE TEPER‡

February 25, 2016

Abstract. We take a decision theoretic approach to predictive inference. We construct a simple dynamic setup in the presence of inherent uncertainty. At any given time period the decision maker updates her posterior regarding the uncertainty related to the subsequent period. The posteriors reflect the decision maker's preferences, period by period. We study the evolution of the agent's posteriors and provide axioms under which the decision maker exhibits learning in a Bayesian fashion. We show how behavioral implications of different Bayesian models differ from one another, and specifically from those dictated by exchangeability.

Keywords: Learning, exchangeability, Bayesianism, local consistency, dynamic decision problem, almost exchangeable.

JEL Classification: D81, D83

The authors would like to thank Peter Klibanoff, Fabio Maccheroni, Evan Piermont, Kyoungwon Seo and Marciano Siniscalchi for comments. Special acknowledgements are given to Eran Shmaya for discussions which contributed to the paper in its early stages, and to Eddie Dekel for comments that significantly improved the exposition.

†Tel Aviv University, Tel Aviv 69978, Israel and INSEAD, Bd. de Constance, 77305 Fontainebleau Cedex, France. e-mail: lehrer@post.tau.ac.il. Lehrer acknowledges the support of the Israel Science Foundation, Grant 963/15.

‡Department of Economics, University of Pittsburgh. e-mail: rteper@pitt.edu.

1. Introduction

Uncertainty regarding payoff relevant factors prevails many economic models. In dynamic environments of this kind, as time goes by agents can gather information that allows them to update their perception of that uncertainty. When an agent is "properly" updating her perception, she gets to extract the maximal information from the data available and information becomes increasingly useless as time passes. We refer to this phenomenon as *learning*. In this paper we investigate the extent to which a decision maker (henceforth, DM) is learning and the extent to which her learning can be identified solely on the basis of her decisions.

We take a decision theoretic approach to predictive inference as in de Finetti [6, 7] and Doob [8]. We introduce a simple dynamic setup with inherent uncertainty and study the agent's preferences over uncertain alternatives, period by period. We then provide axioms under which an agent learns in a Bayesian fashion. In this setup we show how behavioral implications of various Bayesian models may differ from one another, and specifically from those dictated by exchangeability (de Finetti [6]).

1.1. **Bayesianism and learning.**

Consider a dynamic setup where prior to every period the DM faces uncertainty regarding the upcoming state, and specifies her preferences over uncertain alternatives (Anscombe-Aumman [2] acts) for the following period.[1] Then, during the period itself, while the state is realized and uncertainty is resolved, the DM has to make up her mind regarding the next period.

We assume the DM is a subjective expected-utility maximizer, and thus, following every history of realizations, her preferences are induced by some *posterior* belief over the state space. The aim of this study is to investigate the evolution of such posterior beliefs and to find conditions under which different forms of Bayesian learning take place.

We introduce an essential ingredient of Bayesianism and refer to it as *local consistency*. Consider two alternatives, f and g, available to the DM after every history of past realized states. Following a particular history, she specifies her preferences over

[1]Note that following every history of realizations, preferences are over outcomes determined by one-shot resolution of uncertainty.

f and g. From the point of view of that same history, the DM also conveys her preferences regarding f and g looking one period ahead, without knowing which state is going to be realized in the meantime. Local consistency dictates that following any history the DM prefers f over g if and only if, she prefers f over g looking one period later, assuming she does not know which state has been realized in the meanwhile. Throughout the paper we are going to postulate local consistency.

Clearly, local consistency imposes certain restrictions over the evolution of preferences, and thus on the posteriors. Tomorrow's posteriors have to aggregate in a manner that alternatives are evaluated tomorrow as they are evaluated today, using today's posteriors. In other words, following every history the DM is updating her posterior in a Bayesian fashion.[2] It implies that the posteriors form a martingale.[3]

By representing the Bayesian DM's posteriors as a martingale, we immediately obtain from Doob's martingale convergence theorem that consistency induces learning: the effect of additional information becomes increasingly negligible as the history observed by the DM becomes longer. Moreover, the posteriors converge.

1.2. **Exchangeability.** Local consistency on its own is equivalent to Bayesianism, which, in principle, does not impose any further restrictions on how the DM updates her belief. Exchangeability (de Finetti [6]) is a particular form of Bayesianism that introduces further restrictions over such updating. Suppose, for instance, that there are only two states, H and T. The DM believes that nature tosses a coin in each period in order to determine the state for that period, but she does not know the parameter of the coin; she has only a prior belief over the parameters. Every period, she observes the realized state, updates her belief regarding the real parameter, and thereby her belief regarding the probability over the state in the subsequent period. This scenario is akin to exchangeability: the DM has a broad picture regarding the evolution of states, in light of which she updates her beliefs regarding the outcome of the toss in the following period. From her broad picture one can easily determine,

[2]By Bayesian we mean that following every history, she behaves as if there is a signal that stochastically informs her what the state in the next period will be and allows her to update her belief accordingly.

[3]A set of history-dependent posteriors form a martingale if at any history h, the expected one-period-ahead posterior, equals to the posterior associate with h, where the expectation is taken with respect to the posterior associated with h. See Definition 2 for a formal statement.

for instance, that the empirical frequency of states must converge. Furthermore, the beliefs regarding subsequent states become closer and closer to the empirical frequency.

A Bayesian DM, in contrast, does not necessarily have a broad picture. To illustrate this point, suppose that after a particular history the DM believes the chance of H is p. Furthermore, assume that she believes that if H is realized then in the subsequent period the chance of H increases to $\frac{(1+p)}{2}$, while if T is realized in the next period, the chance of H decreases to $\frac{p}{2}$. This is the rule the DM follows when she updates her posterior. What makes this DM Bayesian (and therefore, time consistent) is that her expected next period posterior (i.e., $p\frac{(1+p)}{2} + (1-p)\frac{p}{2}$) equals to her present posterior (i.e., p). When the DM updates her posterior, it is sufficient for her to know the updating rule. Obviously, she can infer what will be her posterior following every finite sequence of eventualities, but this requires an immense cognitive capacity. For instance, it is not easy to determine what will be the long run empirical frequency of states. This part of the picture, for example, is missing to the DM. Thus, although the DM can update her posterior step by step, it is based on a local rule, from which it is not easy to extract the grand picture of the process as a whole.

There are several reasons why studying exchangeable processes in the present context is interesting. First, we show that local consistency does not only imply Bayesianism, but implies a specific form of learning. It turns out, that eventually, the learning pattern dictated solely by local consistency resembles exchangeability: from some period onwards, the learning process looks similar to the learning in an exchangeable process. More formally, there exists an exchangeable process such that with high probability, following any sufficiently long history, the posterior representing the DM's preferences and the posterior associated with the exchangeable process are close to each other. From some point in time, it may seem to an outside observer that the DM very much follows an exchangeable process and not just Bayesian updating.

The second reason exchangeability is interesting in the current context is the simplification it introduces to the updating process. In general, in order to properly update her belief, the DM has to know the probability of every specific history of any length. This requires high cognitive effort. However, when she follows an exchangeable process, she can use a shortcut and as a result, does not have to keep track of a large number of posteriors. In exchangeable processes, the posteriors are determined only by the frequency of states in every history. In other words, two histories that share the same frequency yield the same posteriors. Thus, updating in an exchangeable process requires much less computational and cognitive capacities. The final reason is that exchangeability is frequently applied in the literature.

These points raise a natural question: what is the difference in behavior between a Bayesian that simply exhibits local consistency and a Bayesian whose posteriors generate an exchangeable process? It turns out that the key behavioral difference is related to the one discussed above, namely that preferences that *follow* two histories whose empirical frequencies coincide, must be equal. We refer to this property as *frequency dependence*. Under the expected utility paradigm, frequency dependence implies that the posteriors following histories with identical frequencies are the same. In an exchangeable process, the *ex-ante* probability of histories with identical frequencies must be the same. Here, in contrast, frequency dependence entails that the *posteriors* following two histories with the same frequencies must coincide. It turns out that the two requirements are not equivalent. The fact that the posteriors following histories that have the same frequencies must coincide, does not imply exchangeability without any further assumptions. Example 3 shows that frequency dependence per-se is not strong enough to guarantee exchangeability.

In the case of two states matters are relatively simple. We show that in this case, frequency dependence, along with local consistency, axiomatize exchangeability. However, when more than two states are involved, an additional consistency assumption (the formal discussion of which is deferred to the relevant section) is needed in order to characterize exchangeability.

To conclude this discussion, we would like to emphasize that in this paper we study updating in a local sense. The posteriors (and their dynamics) represent the DM's preferences (and their evolution) over outcomes determined completely by the one-period-ahead state realization. The posteriors naturally induce a *probability* over all infinite histories. The probability of a state following some history is the one assigned to that state by the history-dependent posterior. Such a probability is not directly revealed through the DM's preferences for bets over infinite state realizations; it is merely an auxiliary mathematical entity on which we can rely so as to make succinct statements regarding primitives. This is in contrast to the global approach taken in the learning literature (e.g., Epstein and Seo [11], Klibanoff, Mukerji and Seo [20] and others) and the updating literature following Savage's [24] $P4$ axiom (see Epstein and

Schneider [9], Sinischalchi [26] and Hanany and Klibanoff [16, 17]). In these papers, the probability explicitly represents the primitives. The local approach introduced here gives rise to a different notion of Bayesianism than in the existing literature, which raises intriguing questions that call for new techniques (see also Fortini, Ladelli and Regazzini [12]).

1.3. **The structure of the paper.** In the subsequent section we discuss the model and the primitives. Section 3 presents the local consistency axiom and shows the ways in which it is related to Bayesianism and learning. Section 4 discusses the frequency-dependence axiom and exchangeability. Section 5 presents the additional consistency axiom and the general exchangeability axiomatization result. The related literature and additional comments are discussed in Section 6. Lastly, all the proofs appear in the Appendix.

2. The Model

2.1. **Acts.** We fix a finite *state space* S. An *act* is a real valued function defined over S that takes values between 0 and 1. An act is interpreted as the utility the decision maker (DM) derives, depending on the realized state.[4] Denote the set of acts by $\mathcal{A}$. The classical theory of decision making under uncertainty[5] deals with preferences $\succeq$ over $\mathcal{A}$.

In a framework similar to the one described above, Anscombe and Aumann [2] put forth the foundation for *Subjective Expected Utlity (SEU)* theory. A possible interpretation of this theory is that the decision maker entertains a prior probability[6] $p \in \Delta(S)$ over the states S where an act $f \in \mathcal{A}$ is being evaluated according to its expected value $\mathbb{E}_p(f) = \sum_{s \in S} p(s) f(s)$. We assume throughout that *every* preference

[4] One can also consider the classical Anscombe-Aumann [2] set-up. In this case, standard axioms imply that the vNM utility index can be identified and that the formulation of alternatives as utility acts, as we use here, is well defined. Such results, on which we rely, have been established repeatedly.

[5] See, for example, Savage [24], Anscombe and Aumann [2], Schmeidler [25], Gilboa and Schmeidler [13], Karni [18], Bewely [3], Klibanoff, Marinacci, and Mukerji [19], Maccheroni, Marinacci and Rustichini [23], Chateauneuf and Faro [5], Lehrer and Teper [22], and many others.

[6] For a finite set A, denote by $\Delta(A)$ all probability distributions over A.

relation discussed satisfies the Anscombe-Aumann assumptions and admits an SEU representation.

2.2. **Dynamics.** Consider a scenario whereby every period a state of nature is materialized and the DM derives utility that depends on the realized state and the chosen alternative. We assume that alongside with the utility derived, as the history of realizations evolves, the DM might change her beliefs regarding the likelihood of future events. As a result, her preferences might change as well.

A sequence $(s_1, ..., s_k) \in S^k$ is said to be a *history of states* of length k. The set of all possible histories is $\mathcal{H} := \cup_{k=0}^{\infty} S^k$. A typical history will be denoted by h. We assume that a DM is characterized by a collection of preferences over $\mathcal{A}$, indexed by histories h. Formally, for every history $h \in \mathcal{H}$ there is a preference $\succeq_h$ defined over $\mathcal{A}$ and interpreted as the DM's preferences following the history h.

We assume throughout that every history-dependent preferences $\succeq_h$ admit an SEU representation with respect to the probabilities p_h. That is, after every history, say h, the decision maker entertains some *posterior* over the state space S, that might depend on h, and evaluates acts in $\mathcal{A}$ according to their expected value.

The main question asked in this paper is what are the additional properties needed that tie together the different history-dependent preferences, and allows to determine that the DM is learning, that is, that the posteriors, p_h, are evolving according to some learning procedure as the histories evolve.

2.3. **Infinite histories and the generated probability.** Let $\Omega = S^{\mathbb{N}}$ be all possible infinite histories. A finite history $h \in S^k$ can also be considered a subset of Ω by looking at all those infinite histories such that their projection to the first k periods is exactly h. Thus, if a history $h' \in S^n$ is a continuation of $h \in S^k$ (that is, $n \geq k$ and the projection of h' into the initial k periods is exactly h), then we will use the notation $h' \subseteq h$. Similarly, if an infinite history $\omega \in \Omega$ is a continuation of $h \in S^k$ (that is, the projection of ω into the first k periods is exactly h), then we will use the notation $\omega \in h$.

Consider the set of infinite histories Ω endowed with the σ-algebra generated by the finite histories $\mathcal{H}$. Denote by μ the unique countably-additive probability over Ω that is *generated* by the posteriors $\{p_h\}_{h \in \mathcal{H}}$: for every μ-positive history $h \in \mathcal{H}$, the μ-probability that s follows h is $p_h(s)$. If no confusion might occur, we will refer to μ as the probability (over the infinite histories tree). In certain cases it would be convenient to think about the evolution of the posteriors $\{p_h\}_{h\in\mathcal{H}}$ by means of the probability μ they generate.

Remark 1. *Bayesianism in the context of decision theory has been discussed broadly in the literature. Following Savage's* [24] *axiom P4, the relation between an axiom (or a class of axioms) referred to as* dynamic consistency *and Bayesianism has been investigated (e.g., Hammond* [14, 15]*, Chapter 7 in Tallon and Vergnaud* [27] *and the references within). Dynamic consistency as discussed in that literature is conceptually ill-suited for our setup, as it would typically apply to acts over the "global" state space* Ω*. In contrast, we consider "local behavior", and "beliefs" that, given each history, correspond to preferences over choice objects defined over* S*, the stage-state space.*

The measure μ *defined over subsets of* Ω *is merely a mathematical entity which does not directly represent the DM's preferences.*[7] *We introduce it in order to ease notation, while all the results in this paper can be stated in terms of the parameters representing the primitives. It is important to note that the construction of the probability* μ *places no further assumptions on preferences beyond expected utility.*

In particular, this hints to a conceptual difference between global Bayesianism and local Bayesianism; global analysis does not guarantee any type of local Bayesianism and learning as discussed here. For instance, one could reinterpret P4 (or, dynamic consistency) to fit our setup. It turns out that it would not guarantee that posteriors *evolve according to Bayesian updating. This will be evident from our axiomatization of local Bayesianism as formulated in Proposition 1. See also Remark 2 below.*

2.4. **Exchangeability and Bayesianism.** As discussed in the Introduction, in this paper we strive to focus on general Bayesianism and exchangeability. Our first step is to find conditions under which an agent, exhibiting preferences as those described in the framework above, is Bayesian.

Consider a DM who believes that a distribution (or, a parameter) p over the state space S is chosen randomly according to some distribution θ. In this case, states are selected according to the chosen p in an i.i.d. manner. The DM is not aware of the data-generating parameter; She has nothing but a belief regarding the way in which this parameter is chosen. At every given period she observes the realized state and

[7] The primitive is not the DM's ranking of contingent outcomes depending on the realization of Ω.

based on the history of the realized states she attempts to learn the identity of the real parameter. After each history h she Bayesian updates her belief (initially θ) to obtain θ_h. She then calculates the expected posterior $\mathbb{E}_{\theta_h}(p)$ in order to make decisions. It turns out that as the number of observations increases, the DM's belief converges to the true "parameter" that governs the process.

Let $h \in \mathcal{H}$ be a history. Denote by $\phi_h(s)$ the frequency of state s in h –the number of times s occurred during the history h. When two histories h and h' satisfy $\phi_h = \phi_{h'}$, we say that they *share the same frequency*. Notice that if h and h' share the same frequency it follows that they are of the same length.[8] The following definition is due to de Finetti [6]:

Definition 1. *μ is* exchangeable *if whenever two histories $h, h' \in \mathcal{H}$ share the same frequency, $\mu(h) = \mu(h')$.*

de Finetti showed that exchangeable processes are interesting because they are characterized by i.i.d. conditionals, as described above. Kreps (1988) summarizes the importance of this result:

> "*...de Finetti's theorem, which is, in my opinion, the fundamental theorem of statistical inference – the theorem that from a subjectivist point of view makes sense out of most statistical procedures.*"

While in exchangeable processes the DM obtains signals that are i.i.d. conditional on the state-generating parameter, a general Bayesian considers an abstract signal structure conditional on the true parameter. We provide a consistency condition that will guarantee that posteriors correspond to Bayesian updating. We then move on to study exchangeable processes, and provide an additional condition that guarantees not only that the DM is Bayesian, but also that her posteriors follow a learning pattern akin to exchangeable process.

3. Consistency

Let h be a history of length k and s a state. By hs we denote the history of length $k+1$, starting with h and ending with s.

[8] Note that h and h' share the same frequency, not the same relative frequency.

For state $s \in S$, act $f \in \mathcal{A}$ and utility $a \in [0,1]$, $f_{-s}a$ stands for the act that yields utility $f(s')$ for every state $s' \neq s$ and utility a in state s. A state $s \in S$ is $\succeq_h$-*null* if for every act $f \in \mathcal{A}$ and utility $a \in [0,1], f \sim_h f_{-s}a$. A history $h' = hss_1...s_k \in \mathcal{H}$ is *null* if s is $\succeq_h$-*null*. It is clear that every subsequent history of a null history is also null. Since we assume preferences adhere to SEU, a $\succeq_h$-null state is one whose p_h-probability is 0. Thus, the history h' is null if its μ-probability is 0.

Let $\mathcal{A}_c$ be the collection of all constant acts. That is, if $f \in \mathcal{A}_c$ then $f(s) = f(s')$ for all $s, s' \in S$. Fix an act $f \in \mathcal{A}$, a history h and a state $s \in S$. Denote by $c_{hs}(f)$ the constant equivalent (in $\mathcal{A}_c$) of f after the history hs; that is, $c_{hs}(f) \sim_{hs} f$. Now, define an act $\hat{c}_h(f) \in \mathcal{A}$ by $\hat{c}_h(f) = (c_{hs_1}(f), ..., c_{hs_{|S|}}(f))$. That is, in case s occurs, the reward under the act $\hat{c}_h(f)$ is $c_{hs}(f)$. The act $\hat{c}_h(f)$ represents the way in which the DM perceives f, looking one period ahead into the future. The following axiom postulates that, given a non-null history h, the DM is indifferent between f and the way she perceives it looking one period ahead into the future, without knowing which state will be realized.

Local Consistency: For every non-null history h and act f,

$$f \sim_h \hat{c}_h(f).$$

Example 1. *Suppose that $S = \{H, T\}$. Suppose also that at the beginning of the sequential decision problem the DM believes that H and T are being determined by a toss of a fair coin. If at the first period the realized state is H, she then "updates" her belief and from then on believes that H and T will be determined by an infinite toss of a $\frac{2}{3}$-biased coin. But if the realized state is T, she comes to believe that H and T will be determined (from that point on) by an infinite toss of a $\frac{1}{3}$-biased coin. Assume that $\succeq_h$ reflects SEU maximization with respect to the beliefs just described.*

In other words, $p_\emptyset = (\frac{1}{2}, \frac{1}{2})$, where $\emptyset$ stands for the empty history (at the beginning of the process). Consequently, for every history h starting with H, $p_h = (\frac{2}{3}, \frac{1}{3})$, and for every history h starting with T, $p_h = (\frac{1}{3}, \frac{2}{3})$. The decision making process based on these posteriors demonstrates local consistency. *Nevertheless, it is not exchangeable since the probability of HTT is different than that of THT.*

Definition 2. *The posteriors* $\{p_h\}_{h\in\mathcal{H}}$ *form a* martingale *if for every history* $h \in \mathcal{H}$*,*

$$\mathbb{E}_{p_h}(p_{hs}) := \sum_{s\in S} p_h(s)p_{hs} = p_h.$$

It is well known that the dynamics of posteriors associated with Bayesian updating is characterized by a martingale. This facilitates the next axiomatization result, stating that *local consistency* is necessary and sufficient for Bayesianism.

Proposition 1. Local consistency *is satisfied if and only if the posteriors form a martingale.*

Remark 2. *A different formulation of consistency would require that for every two acts* f *and* g *and every history* h*, if* $f \succeq_{hs} g$ *for any* $s \in S$*, then* $f \succeq_h g$*. This could be seen as a reinterpretation of Savage's* [24] *P*4 *(or, dynamic consistency) to our setup. However, while this property implies that the posterior* p_h *is an average of the posteriors* $\{p_{hs}\}_{s\in S}$*, it does not guarantee they form a martingale. Moreover, if we were to adopt the standard approach in which objects of choice following every history are sequential, where the standard formulation of dynamic consistency can be postulated, we would still obtain that the posterior following any history is an average of the one period ahead posteriors, but not necessarily a martingale. Thus, local Bayesianism would not be guaranteed.*

*This is not merely a technical detail, but stems from a conceptual difference between global Bayesianism and local Bayesinism. In the global approach, objects of choice are sequential and following every history are a function of the residual uncertainty. In this case, P*4 *and its reformulations guarantee that before and after a state is realized, the DM assesses the probability of the residual uncertainty in an identical manner. On the other hand, in our setup the DM is not asked to provide probabilities of residual uncertainty, but the probabilities of events recurring given information regarding the history. These probabilities need not be the same before and after the realization of uncertainty in some period of time. Local consistency then puts restrictions as to how these posterior probabilities can change upon such realizations, and guarantees updating in a Bayesian fashion.*

Definition 3. *The probability μ is* almost exchangeable *if there exists an exchangeable probability ζ such that for every $\epsilon > 0$ there exists $T > 0$ such that*

$$|\mu(hs|h) - \zeta(hs|h)| < \epsilon \tag{1}$$

for a set of histories $h \in S^T$ whose probability is at least $1 - \epsilon$, and for every $s \in S$.

A process μ is almost exchangeable when we can construct an exchangeable process ζ such that, as time goes by, for every history in a set of histories of μ-probability close to 1, the posteriors associated with the two processes following that history are close. In particular, while early on updating is not necessarily similar to that of an exchangeable process, at some point in time and on it is.

The next result states that any Bayesian DM is following learning patterns akin to almost exchangeable process.

Proposition 2. *If the posteriors form a martingale, then the probability is almost exchangeable.*

Combining the results above, we obtain the following theorem:

Theorem 1. *If* local consistency *is satisfied, then the probability is almost exchangeable.*

The significance of the results is in tying together the different threads of the decision making process (following every history) into one coherent theory, and providing a behavioral foundation for Bayesianism and learning in the current setup. *Local consistency* tells us that the DM keeps learning as she observes more and more realizations of states. Without *local consistency*, the beliefs could be completely unrelated given the histories. The role of the axiom is to provide a consistent manner in which the beliefs evolve. The DM updates her belief in a Bayesian fashion as she observes past realizations, and as time goes by, her belief converges. This implies that in the long run getting additional information will not change significantly the forecast of nearby events. Furthermore, as time goes by the belief concerning future evolution of states becomes less complex: it converges to an i.i.d. process (i.e., it is determined solely by the one-stage distributions).

As discussed in Remark 1, while the probability μ entails information regarding the DM preferences, it is constructed by us modelers and does not directly represent the DM's beliefs over Ω. Theorem 1 (and subsequent results) can be interpreted as follows: if *local consistency* is satisfied then the modeler knows the DM maker is going to learn in an almost exchangeable fashion, where learning is accomplished with respect to the natural measure (μ), which is consistent with the DM's revealed beliefs. In a similar manner, the modeler can appeal to the DM's reasoning. If she satisfies *local consistency*, then she is guaranteed to learn in an almost exchangeable manner with respect to the natural measure induced by her beliefs.

Remark 3. *By assuming a stronger version of* local consistency*, in which $f \sim_h \hat{c}_h(f)$ for* every *history h and act f, it is possible to reformulate Theorem 1 to incorporate learning as in Definition 3 even when the DM is "surprised" and a null history actually occurs. Ex-ante such histories are of probability 0, but, in the interim, even in case that such a history (h) is realized, and the stronger* local consistency *is satisfied, the probability generated by the posteriors associated with the continuation histories (of history h) is almost exchangeable.*[9] *Similar modifications can be made to subsequent theorems as well.*

4. Frequency-Determined Preferences

The process in Example 1 satisfies *local consistency*, and is therefore almost exchangeable, but not exchangeable. We now ask what property exchangeable processes posses that almost exchangeable processes do not.

Frequency Dependence: For every two non-null histories h and h' such that $\phi_h = \phi_{h'}$,

$$\succeq_h = \succeq_{h'} .$$

Frequency dependence postulates that preferences associated with (positive μ-probability) histories that share the same frequency are identical. Note that, given our assumption

[9]This construction is similar to the one behind *Lexicographic Probability Systems* presented in Brandenburger, Friedenberg and Keisler [4] (see Definition 4.1), where the probabilities in the system are almost exchangeable. One difference is that there is no (linear) ordering of the probabilities, but a partial ordering naturally inherited from the partial ordering of histories.

that preferences following every history adhere to SEU, and given the uniqueness of the representation of SEU preferences, *frequency dependence* implies that $p_h = p_{h'}$.

Example 2. *An urn contains two biased coins: A and B. One coin is randomly selected and then flipped infinitely many times. Suppose that the probability of Heads (H) coming up when coin A is tossed is* $\frac{1}{3}$*, whereas the probability is* $\frac{2}{3}$ *when coin B is tossed. Ex-ante, the probability that the outcome would be H or T in any given period is* $\frac{1}{2}$ *.*

Let $S = \{H, T\}$ *and assume that the DM believes states are selected according to the description above. At the beginning of the process, the belief regarding the identity of the selected coin was uniform* ($\frac{1}{2}$ *on each coin). After observing a history h, the DM updates her belief, and based on the updated belief she then takes a decision. Given the nature of the process described, the posterior probability the DM assigns to the coins after two histories that share the same frequency is the same, implying that* frequency dependence *is satisfied.*

This process serves also a classic example of an exchangeable process. As the histories observed by the DM grow longer and longer, she updates her belief and the identity of the selected coin grows clearer and clearer. If coin A was selected initially, then the relative frequency of H's converges with time to $\frac{1}{3}$ *(with probability 1, of course), while if coin B was selected, the relative frequency of H's converges with time to* $\frac{2}{3}$*. After a while the DM faces a process that is close to an i.i.d. process. As illustrated in this example, every exchangeable process possesses* frequency dependence.

Frequency dependence assumes that *posteriors* following two histories sharing the same frequency coincide, while in an exchangeable process, two histories that share the same frequency have the same *probability.* Histories that have the same frequency are used in both concepts. A DM who believes that nature selects states according to an exchangeable process and maximizes SEU satisfies *frequency dependence*. The converse, however, is not true, which makes the next result all the more interesting and challenging. In order to convince the reader, consider the following example.

Example 3. *Let* $S = \{H, T\}$*. The first two periods of the process are as in Example 2: at the first period the states from S are chosen with equal probabilities. At the second period a* $\frac{2}{3}$ *or* $\frac{1}{3}$*-biased coin is tossed according to whether H or T was realized in the first period. From the third period onwards the process is i.i.d: following the history*

HH the continuation is forever H and following TT the continuation is forever T. Finally, if the history is mixed (either HT or TH) the process goes on according to a toss of a fair coin.

It is clear that after any two histories (that occur with positive probability) that share the same frequency, the posteriors coincide. Thus, frequency dependence *is satisfied. This process, however, is not exchangeable: the probability of HTT is positive while the probability of TTH is 0. Moreover, the underlying process does not satisfy* local consistency*: following T, the probability of H is $\frac{1}{3}$, while the probability of H in the subsequent period is $\frac{1}{3} \cdot \frac{1}{2} + \frac{2}{3} \cdot 0 = \frac{1}{6}$.*

The example shows that even though the process satisfies *frequency dependence*, it is not exchangeable. The following theorem states that this is no longer the case when *local consistency* is satisfied, and that these two properties characterize exchangeable processes when the state space consists of two states.

Theorem 2. *Suppose the state space $S = \{H, T\}$ consists of two states. Then,* local consistency *and* frequency dependence *are satisfied if and only if μ is exchangeable.*

5. A More-than-Two-States Space

The assumption in Theorem 2 that S consists of two states guarantees, without relying on *frequency dependence*, that the probability of HT equals the probability of TH, which is implied by exchangeability. Example 4 below shows that when S consists of more than two states, the probabilities of HT and TH are not necessarily equal even though *local consistency* is satisfied.

Example 4. *Suppose that $S = \{H, T, M\}$. Suppose that the DM believes at the beginning that H,T and M are selected according to $p_\emptyset = (\frac{1}{2}, \frac{1}{4}, \frac{1}{4})$. From the second stage on the process becomes i.i.d. with a distribution that depends on the outcome of the first state. Following H, the one-stage distribution is $p_H = (\frac{1}{2}, 0, \frac{1}{2})$; following T, the one-stage distribution is $p_T = (\frac{1}{3}, \frac{2}{3}, 0)$; and following M it is $p_M = (\frac{2}{3}, \frac{1}{3}, 0)$. A DM acting according to these beliefs would satisfy* local consistency*. Nevertheless, μ generated by this beliefs is such that $\mu(HT) = 0$ while $\mu(TH) > 0$.*

Let $S_1, S_2, ...$ be a stochastic process determining the states in every period and denote by ν the distribution it generates. We say the process is *k-stationary* if for every history h of length k, $\nu((S_1...S_k) = h) = \nu((S_tS_{t+1}...S_{t+k-1}) = h)$ for any time t. When a process is k-stationary for every k we say that it is *stationary.* While *local consistency* implies that the underlying process is 1-stationary, it does not imply that it is stationary. The following example shows that the axioms above do not imply that the underlying process is 2-stationary.

Example 4 continued. *The main point of this example is that $\mu(S_1S_2 = HT) = 0$, while $\mu(S_2S_3 = HT) = \frac{1}{4} \cdot \frac{1}{3} \cdot \frac{2}{3} + \frac{1}{4} \cdot \frac{2}{3} \cdot \frac{1}{3} > 0$. In particular, $\mu(S_1S_2 = HT) \neq \mu(S_2S_3 = HT)$, making the process non-stationary.*

From Example 4 and the discussion above, it seems that in order to be able to prove a representation result for exchangeable processes in a more-than-two-states space, additional structure is called for in terms of *local consistency.* Such an axiom would guarantee that the underlying process is 2-stationary.

By $\mathbb{1}_r$ we denote the act that induces utility 1 if state $r \in S$ is realized and by 0 if otherwise. Fix a history $h \in \mathcal{H}$ and states $s, r \in S$. As in Section 3, $c_{hs}(\mathbb{1}_r)$ is the constant equivalent of $\mathbb{1}_r$ if the history is hs. Now, the act $c_{hs}(\mathbb{1}_r)\mathbb{1}_s$ is the act that induces utility $c_{hs}(\mathbb{1}_r)$ if state s is realized and 0 if otherwise. Let $c_h(s, r)$ be the constant equivalent of $c_{hs}(\mathbb{1}_r)\mathbb{1}_s$ following the history h (that is, $c_h(s, r) \sim_h c_{hs}(\mathbb{1}_r)\mathbb{1}_s$). While notation is a bit cumbersome, $c^1_h(s, r)$ simply represents the DM's valuation (given history h) of placing a bet on state s for today followed by a bet on state r for tomorrow.

The same idea can be repeated for the history hw (where $w \in S$ is some state); that is, $c_{hws}(\mathbb{1}_r)\mathbb{1}_s$ is the act inducing the constant equivalent of $\mathbb{1}_r$ (following the history hws) at state s, and 0 otherwise. Evaluating $c_{hws}(\mathbb{1}_r)\mathbb{1}_s$ from the point of view of history hw yields the constant equivalent $c^1_{hw}(s, r)$. Now, following history h, let $c^2_h(s, r)$ be the constant equivalent of $(c^1_{hw_1}(s, r), ..., c^1_{hw_{|S|}}(s, r))$. That is, $c^2_h(s, r)$ represents the DM's valuation (given history h) of placing a bet on state s for tomorrow, followed by a bet on state r for the day after tomorrow (without knowing or conditioning on today's state realization).

The following axiom captures a idea similar to *local consistency*; we postulate that the DM is indifferent between a bet on s today followed by a bet on r tomorrow (if

s indeed occured), and a bet on s tomorrow followed by a bet on r in the following period (provided that s has indeed occured).

Two-Tier Local Consistency: For every non-null history h and states s and r,

$$c_h^1(s,r) \sim_h c_h^2(s,r).$$

Theorem 3. *Let S be any state space. In this case,* local consistency*,* two-tier local consistency *and* frequency dependence *are satisfied if and only if μ is exchangeable.*

Remark 4. *The preferences in Example 4 satisfy* local consistency *but not* frequency dependence *(both HM and MH are histories with positive probability but their associated posteriors are different). As we have seen, the resulting probability for such preferences is neither 2-stationary nor exchangeable. It is possible that in the presence of* frequency dependence*,* local consistency *implies 2-stationarity (and exchangeability from Theorem 3), yet this question remains open at the current point.*

6. The Literature and Additional Comments

6.1. **Literature.** Several papers have studied different aspects of learning, including exchangeability, in the context of decision theory (e.g., Epstein and Seo [11], Klibanoff, Mukerji and Seo [20] and Al-Najjar and De Castro [1]). These papers are typically not interested in exchangeability in the context of expected utility. The reason is that in the setup these papers consider, the outcome depends on the realizations of all states across time.[10] In such a setup, axiomatizing exchangeability in the context of expected utility is a straightforward application of de Finetti's [6] characterization and is obtained by assuming symmetry of preferences for (finite) permutations of the experiments' outcomes.

Epstein and Schneider [9], Sinischalchi [26] and Hanany and Klibanoff [16, 17] discuss dynamic models of ambiguity and issues that emerge as a result of updating vis-a-vis ambiguity. Specifications concerning the meaning of learning in the long-run is discussed in detail in Epstein and Schneider [10] without an axiomatic foundation.

[10]As opposed to a sequential problem in which every period produces an outcome that depends on the state realized at that period.

In a statistical framework, Fortini, Ladelli and Regazzini [12] characterize a collection of posteriors consistent with an exchangeable process. One of their properties is similar to our *frequency dependence* when translated to posteriors, but this paper is primarily concerned with exchangeable processes. Our focus, in contrast, is the relations and gaps between general Bayesian processes and an exchangeable ones.

6.2. **Symmetry in the current framework.** It is possible to use a notation developed specifically for *two-tier local consistency* and formulate an axiom stating that, from the point of view of the current period, the DM is indifferent between 1. the bet on state r today and then on state s (conditional on r occurring); and 2. the bet on state s today and then on r (conditional on s occurring). This would be a simple symmetry condition for permutations of two possible experiments. It is possible to further develop this idea and formulate, at a high modelling intractability cost, a symmetry condition for every finite permutation of the experiments' outcome. Following de Finetti's characterization, such an axiom will characterize exchangeable processes without any further postulations.

Our consistency axioms require the DM looks forward either one or two periods ahead into the future (depending on the axiom), and makes predictions about the likelihood of subsequent events. Making such predictions would also be a feature of any symmetry axiom we formulate, the main difference being that a fully fledged symmetry axiom would require making predictions while looking forward far into the future (for any finite but unbounded number of periods).[11] This seems a much more complicated task.

6.3. **A probabilistic issue.** It is clear that in any exchangeable process, the posteriors that follow any two (positive probability) histories having the same frequency coincide. As shown in Example 3, however, the inverse direction is typically incorrect.

The local consistency axioms (i.e., *local consistency* and *two-tier local consistency*) imply 1 and 2-stationarity. But from Proposition 1, *time consistency* alone has actually a stronger consequence. It implies that the one-stage predictions form a martingale. As shown, this fact, alongside with *frequency dependence* and *two-tier local consistency*,

[11]Note that making predictions far into the future is different from *frequency dependence* which imposes restrictions on the DM's predictions for the *current period*.

is sufficient to guarantee exchangeability, and in particular stationary. The inverse, however, is incorrect: there could be a case (for instance, in a Markov chain when the initial distribution over states is invariant) in which the underlying process is stationary while the one-stage predictions do not form a martingale. Typically, in such a case, the stage-predictions do not converge, and no learning is taking place.

This observation naturally raises the following questions, which seem to be rather difficult to answer:

1. Is a stationary stochastic process necessarily be exchangeable whenever any two posteriors that follow histories sharing the same frequency coincide .
2. In case the answer to the above question is on the affirmative, what behavioral axiomatization would capture a DM who is a present-value expected utility maximizer when the underlying state of nature evolves according to a stationary process.

References

[1] Al-Najjar, N., and L. De Castro (2014) Parametric representation of preferences, *Journal of Economic Theory*, **150**, 642–667.

[2] Anscombe, F.J., and R.J. Aumann (1963) A Definition of Subjective Probability, *The Annals of Mathematical Statistics*, **34**, 199–295.

[3] Bewley, T.F. (2002) Knightian Decision Theory. Part 1, *Decisions in Economics and Finance*, **25**, 79–110.

[4] Branderburger, A., A. Friedenberg and H.J. Keisler (2008) Admissibility in Games, *Econometrica*, **76**, 307–352.

[5] Chateauneuf, A., and J.H. Faro (2009) Ambiguity Through Confidence Functions, *Journal of Mathematical Economics*, **45**, 535–558.

[6] de Finetti, B. (1930) Fuzione Caratteristica di un Fenomeno Aleatorio, *Atti della Academia Nazionale dei Lincei Rendiconti, Class di Scienze Fisiche, Mathematiche e Naturali*, **4**, 86–133.

[7] de Finetti, B. (1937) La Prvision: ses lois logiques, ses sources subjectives, *Annales de l'Institut Henri Poincarè*, **7**, 1–68.

[8] Doob, J. L. (1953) Stochastic Processes, Wiley, New York.

[9] Epstein, L. G., and M. Schneider (2003) Recursive Multiple Priors, *Journal of Economic Theory*, **113**, 1–31.

[10] Epstein, L. G., and M. Schneider (2007) Learning Under Ambiguity, *Review of Economic Studies*, **74**, 1275–1303.

[11] Epstein, L.G., and K. Seo (2010) Symmetry of Evidence without Evidence of Symmetry, *Theoretical Economics*, **5**, 313–368.

[12] Fortini, S., Ladelli, L., and E. Regazzini (2000) Exchangeability, Predictive Distributions and Parametric Models, *The Indian Journal of Statistics*, **20**, 86–109.

[13] Gilboa, I., and D. Schmeidler (1989) Maxmin Expected Utility with Non-Unique Prior, *Journal of Mathematical Economics* **18**, 141–153.

[14] Hammond, P. (1988) Consequentialist Foundations for Expected Utility, *Theory and Decision*, **25**, 25–78.

[15] Hammond, P. (1989) Consistent plans, Consequentialism, and Expected Utility, *Econometrica*, **57**, 1445–1449.

[16] Hanany, E., and P. Klibanoff (2007) Updating Preferences with Multiple Priors, *Theoretical Economics*, **2**, 261–298.

[17] Hanany, E., and P. Klibanoff (2009) Updating Ambiguity Averse Preferences, *The BE Journal of Theoretical Economics*, **9**, 1–51.

[18] Karni, A. (1993) A Definition of Subjective Probabilities with State Dependent Preferences, *Econometrica*, **61**, 187–198.

[19] Klibanoff, P., M. Marinacci, and S. Mukerji (2005) A Smooth Model of Decision Making Under Ambiguity, *Econometrica*, **73**, 1849–1892.

[20] Klibanoff, P., S. Mukerji, and K. Seo (2014) Perceived Ambiguity and Relevant Measures, *Econometrica*, **82**, 1945 -1978.

[21] Kreps, D. M. (1988) Notes on the Theory of Choice, Westview Press.

[22] Lehrer, E., and R. Teper (2011) Justifiable Preferences, *Journal of Economic Theory*, **146**, 762–774.

[23] Maccheroni, F., M. Marinacci, and A. Rustichini (2006) Ambiguity Aversion, Robustness, and the Variational Representation of Preferences, *Econometrica*, **74**, 1447–1498.

[24] Savage, L. J. (1954) The Foundations of Statistics, Wiley, New York.

[25] Schmeidler, D. (1989) Subjective Probabilities and Expected Utility Without Additivity, *Econometrica,* **57**, 571–587.

[26] Sinisalchi, M. (2011) Dynamic Choice Under Ambiguity, *Theoretical Economics*, **6**, 379–421.

[27] Tallon, J. M., and J. C. Vergnaud (2006) Knowledge, Beliefs and Economics, R. Arena and A. Festre eds, Edward Elgar publishers.

Appendix A. Proofs

A.1. **Proof of Proposition 1.** In order to see the sufficiency of *local consistency*, fix a history h and let $\mathbb{1}_E$ denote the indicator function for an event $E \subseteq S$. We know that from *local consistency* it implies that $\mathbb{1}_E \sim_h \hat{c}_h(\mathbb{1}_E)$. Since $\succeq_h$ is SEU represented by p_h, we have that $p_h(E) = \sum_{s\in S} p_h(s)c_{hs}(\mathbb{1}_E)$. By the definition of c_{hs}, we have that the right hand side of the latter equality equals $\sum_{s\in S} p_h(s)p_{hs}(E)$. Since E is arbitrary and since the mixture weights, $p_h(s)$, do not depend on E, we have that $p_h = \sum_{s\in S} p_h(s)p_{hs}$, implying that the posteriors form a martingale.

The necessity of *local consistency* is immediate and is thus omitted.

A.2. **Proof of Theorem 1.** For every $\omega = (\omega_1, \omega_2, \dots) \in \Omega$ and T, denote by $\omega^T = (\omega_1, \dots, \omega_T)$ the prefix of ω whose length is T, and for $T_1 < T_2$ denote $\omega^{T_1,T_2} = (\omega_{T_1}, \dots, \omega_{T_2})$.

Assume *local consistency*, then by Proposition 1 the posteriors form a martingale. Thus, the posteriors converge μ-almost surely (see Doob [8]). Let $C \subseteq \Omega$ be a set of μ-probability 1 for which the martingale of posteriors converges. That is, for every $\omega \in C$, $p_{_{\omega^T}}$ converges (as $T \to \infty$) to a limit denoted as p_ω. This implies that for every $\varepsilon > 0$ there is $t_1(\varepsilon/3)$ large enough so that[12] for every $T \geq t_1(\varepsilon/3)$

$$\mu\left(\omega \in C;\ ||p_\omega - p_{_{\omega^T}}|| < \varepsilon/3\right) > 1 - \varepsilon/3. \tag{2}$$

Without loss of generality we may assume that for every $\omega \in \Omega$ and T, ω^T has a μ-positive probability (otherwise we can omit this set of ω's, which is measurable.)

Fix $\varepsilon > 0$ and $\omega \in C$. Consider the (history) space conditional on $\omega^{t_1(\varepsilon/3)}$: we examine the frequency of states along sufficiently long continuations of $\omega^{t_1(\varepsilon/3)}$. By the strong law of large numbers there is $t_2(\varepsilon/3)$ such that with probability of at least $1-\varepsilon/3$ (conditional on $\omega_1^t(\varepsilon/3)$), for any $t_2 \geq t_2(\varepsilon/3)$ the frequency of $s \in S$ in $\omega^{t_1(\varepsilon/3),t_2}$ is $\varepsilon/3$ far from the average posterior of s along $\omega^{t_1(\varepsilon/3),t_2}$. Formally, for any $\varepsilon > 0$, there is $t_2(\varepsilon/3) = t_2(\varepsilon/3, \omega)$ such that for any $t_2 \geq t_2(\varepsilon/3)$ and every $s \in S$,

$$\mu\left(\omega \in \Omega;\ \frac{|\sum_{t=t_1(\varepsilon/3)}^{t_2}[\mathbb{1}_{\omega_t=s} - p_{_{\omega^t}}(s)]|}{t_2 - t_1(\varepsilon/3) + 1} < \varepsilon/3 \text{ for every } s \in S \;\middle|\; \omega^{t_1(\varepsilon/3)}\right) > 1 - \varepsilon/3. \tag{3}$$

Recall that for $\omega \in C$, all posteriors, $p_{_{\omega^T}}$, for $T > t_1(\varepsilon/3)$, are close to p_ω up to $\varepsilon/3$. Denote by $\hat{\phi}(h)$ the relative frequency of history h. Due to (2) we obtain,

$$\mu\left(\omega \in C;\ ||\hat{\phi}(\omega^{t_1(\varepsilon/3),t_2}) - p_{_\omega}|| < 2\varepsilon/3\right) > 1 - 2\varepsilon/3. \tag{4}$$

[12]For two probability distributions p, p' over S, we denote $||p - p'|| < \epsilon$ if $|p(s) - p'(s)| < \varepsilon$ for every $s \in S$.

If $t_2(\varepsilon/3)$ is large enough, then $\frac{t_1(\varepsilon/3)}{t_2(\varepsilon/3)} < \varepsilon/3$ and the first $t_1(\varepsilon/3)$ states in $\omega^{t_2(\varepsilon/3)}$ have weight smaller than $\varepsilon/3$. Thus, (4) changes to

$$\mu\left(\omega \in C;\ ||\hat{\phi}(\omega^{t_2}) - p_\omega|| < \varepsilon\right) > 1 - 2\varepsilon/3. \tag{5}$$

In words, the probability of $\omega \in C$ such that the empirical frequency of the states over the history ω^{t_2} (recall, $t_2 \geq t_2(\varepsilon/3)$) is close to p_ω, is at least $1 - 2\varepsilon/3$. Recalling the definition of C, we conclude that with high probability the empirical frequency over histories ω^{t_2} and their posteriors are close to each other. Formally, with the help of (2), when $t_2 \geq t_2(\varepsilon/3) > 3t_1(\varepsilon/3)/\varepsilon$,

$$\mu\left(\omega \in C;\ ||\hat{\phi}(\omega^{t_2}) - p_{_{\omega^{t_2}}}|| < 4\varepsilon/3 \text{ and } ||\hat{\phi}(\omega^{t_2}) - p_\omega|| < \varepsilon\right) > 1 - 2\varepsilon/3. \tag{6}$$

So much for μ.

We turn now to the definition of the exchangeable process, ζ. Recall that p_ω is a random variable defined on C that takes values in $\Delta(S)$. Thus, p_ω and μ induce a probability distribution over $\Delta(S)$, which we denote as θ. Let the parameter space be $(\Delta(S), \theta)$. For every $p \in \Delta(S)$ let B_p be the i.i.d. process with stage-distribution p. ζ is the process defined by the distribution over $\{B_p; p \in \Delta(S)\}$ induced by the distribution θ over the p's (i.e., over $\Delta(S)$).

The definition of ζ and the fact that it is exchangeable imply the following facts:
(i) For ζ-almost every $\omega \in \Omega$, the limit of the empirical frequencies, $\lim_{t\to\infty} \hat{\phi}(\omega^t)$, exists. Denote it by q_ω. q_ω is a limit of empirical frequencies of states in S and is therefore in $\Delta(S)$.
(ii) For ζ-almost every $\omega \in \Omega$, the posteriors w.r.t. to ζ, denoted $p^\zeta_{\omega^t}$, converge to q_ω as $t \to \infty$. That is, $\lim_{t\to\infty} p^\zeta_{\omega^t} = q_\omega$ with ζ-probability 1.
(iii) The distributions of p_ω (induced by μ) and that of q_ω (induced by ζ), both over $\Delta(S)$, coincide.
(iv) For θ-almost every $p \in \Delta(S)$ with ζ high probability, if the relative frequency $\hat{\phi}(\omega^T)$ is close to p, then the posterior (the one induced by ζ following ω^T) must also be close to p. Formally, for θ-almost every $p \in \Delta(S)$ and for every $\varepsilon > 0$, there is t_3 such that for every $t > t_3$,

$$\zeta\left(\omega;\ ||\hat{\phi}(\omega^t) - p|| < \varepsilon/2 \text{ implies } ||p^\zeta_{\omega^t} - p_\omega|| < \varepsilon\right) > 1 - \varepsilon/3. \tag{7}$$

This, in turn, implies that for every $\varepsilon > 0$, there is $t_3(\varepsilon)$ and $\Delta(\varepsilon) \subseteq \Delta(S)$ such that $\theta(\Delta(\varepsilon)) > 1 - \varepsilon/3$ and for every $t > t_3(\varepsilon)$,

$$||\hat{\phi}(\omega^t) - p|| < \varepsilon/2 \text{ implies } ||p^{\zeta}_{\omega^t} - p_\omega|| < \varepsilon. \tag{8}$$

We return to now μ. Due to fact (iii), there is $t_4(\varepsilon)$ such that for every $t_4 > t_4(\varepsilon)$,

$$\mu\Big(\omega;\ \text{there is } p \in \Delta(\varepsilon) \text{ such that } ||\hat{\phi}(\omega^{t_4}) - p|| < \varepsilon/2\Big) > 1 - \varepsilon/3. \tag{9}$$

From (6) and (9) we obtain that for every $t > \max\{t_2(\varepsilon/3), t_4(\varepsilon)\}$,

$$\mu\Big(\omega \in C;\ ||\hat{\phi}(\omega^t) - p_{_{\omega^t}}|| < 4\varepsilon/3, ||\hat{\phi}(\omega^t) - p_\omega|| < \varepsilon \\ \text{and there is } p \in \Delta(\varepsilon) \text{ such that } ||\hat{\phi}(\omega^t) - p|| < \varepsilon/2\Big) > 1 - \varepsilon. \tag{10}$$

Combining with (8) we get that when $t > \max\{t_2(\varepsilon/2), t_3(\varepsilon), t_4(\varepsilon)\}$,

$$\mu\Big(\omega \in C;\ ||\hat{\phi}(\omega^t) - p_{_{\omega^t}}|| < 4\varepsilon/3, ||\hat{\phi}(\omega^t) - p_\omega|| < \varepsilon \\ \text{and } ||p^{\zeta}_{\omega^t} - p_\omega|| < \varepsilon\Big) > 1 - \varepsilon. \tag{11}$$

Thus (due to the triangle inequality),

$$\mu\Big(\omega \in C;\ ||p^{\zeta}_{\omega^t} - p_{_{\omega^t}}|| < 10\varepsilon/3\Big) > 1 - \varepsilon. \tag{12}$$

This completes the proof.

A.3. **Proof of Theorem 2.** It is immediate that the conditions of the theorem are necessary. We prove sufficiency by induction on the length of the history. We start with the base case where $h, h' \in S^2$. We need to show that $\mu(HT) = \mu(TH)$. Denote $p_\emptyset(H) = p, p_H(H) = q$ and $p_T(H) = r$. Then we want to show that

$$p(1 - q) = (1 - p)r. \tag{13}$$

This holds if and only if $p = \frac{r}{1-q+r}$. But this in turn is equivalent to $p = pq + (1-p)r$. The latter, however, holds due to *local consistency*, which means that Eq. (13) holds.

Now assume that the hypothesis holds for any two histories $h, h' \in S^k$ sharing the same frequency, for every history length k smaller than n. We show that $\mu(h) = \mu(h')$ for $h, h' \in S^n$ sharing the same frequency.

Case 1: Assume that $h, h' \in S^n$ share the same frequency and that $h = \bar{h}s$ and $h' = \bar{h}'s$. By definition of μ we have that $\mu(h) = \mu(\bar{h}s) = \mu(\bar{h})p_{\bar{h}}(s)$ and similarly, $\mu(h) = \mu(\bar{h}')p_{\bar{h}'}(s)$. Since h and h' share identical frequencies,

$$\phi(\bar{h}) = \phi(\bar{h}'). \tag{14}$$

From the induction assumption, $\mu(\bar{h}) = \mu(\bar{h}')$. Also implied from Eq. (14) and *frequency dependence* is that $p_{\bar{h}}(s) = p_{\bar{h}'}(s)$. Combined we have that $\mu(h) = \mu(h')$.

Case 2: Assume that $h, h' \in S^n$ share the same frequency and $h = \bar{h}st$ and $h' = \bar{h}'ts$. Then, $\phi(\bar{h}) = \phi(\bar{h}')$, and from the induction assumption $\mu(\bar{h}) = \mu(\bar{h}')$. Relying on *local consistency* and repeating the same arguments as in the base case, one obtains that $\mu(\bar{h}st) = \mu(\bar{h}'st)$.

Case 3: Assume now that $h, h' \in S^n$ share the same frequency and that $h = \bar{h}s$ and $h' = \bar{h}'t$ (where $t \neq s$). This means that t must have been part of the history $\bar{h}$. Let $\hat{h} \in S^{n-2}$ such that $\bar{h} = \hat{h}t$, and similarly, let $\hat{h}' \in S^{n-2}$ such that $\bar{h}' = \hat{h}'s$. We claim that

$$\mu(\bar{h}s) = \mu(\hat{h}ts) \text{ and } \mu(\bar{h}'t) = \mu(\hat{h}'st). \tag{15}$$

If this indeed holds, then $\mu(h) = \mu(\bar{h}s) = \mu(\hat{h}ts)$, where, as we deduce from *Case 2*, the right hand side equals to $\mu(\hat{h}'st) = \mu(\bar{h}'t) = \mu(h')$.

To show that Eq. (15) holds, note that $\phi(\bar{h}) = \phi(\hat{h}t)$ and by the induction assumption $\mu(\bar{h}) = \mu(\hat{h}t)$. Now, similarly to the arguments in *Case 1*, the left equality of Eq. (15) is satisfied. From similar arguments, the right equality holds too.

A.4. **Proof of Theorem 3.** The one thing in the proof of Theorem 2 that does not hold with more than two states and without assuming *two-tier local consistency*, is that $\mu(sr) = \mu(rs)$ for every $s, r \in S$. We prove this point here and by that prove the theorem.

Fix a history $h \in \mathcal{H}$ and states $s, r \in S$. We recursively define $c_h^{n+1}(s, r)$ For $n \geq 3$ by

$$c_h^{n+1}(s,r) \sim_h (c_{hw_1}^n(s,r), ..., c_{hw_{|S|}}^n(s,r)). \tag{16}$$

That is, $c_h^n(s,r)$ reflects the DM evaluation of a bet on state s to occur n periods from now on, followed by the occurance of state r, regardelss of the states realized in the first $n-1$ periods.

Claim 1. *$c_h^1(s,r) \sim_h c_h^n(s,r)$ for every history $h \in \mathcal{H}$ and $n \geq 2$.*

Proof of Claim 1. We prove by induction on n that $c_h^{n-1}(s,r) \sim_h c_h^n(s,r)$ for every history $h \in \mathcal{H}$ and $n \geq 1$. The base case $n = 2$ holds for every history $h \in \mathcal{H}$ due to *two-tier local consistency*. Assume the the statement holds for n. We show it holds for $n+1$. Indeed, by the induction hypothesis $c_{hw}^n(s,r) \sim_{hw} c_{hw}^{n-1}(s,r)$ for every $w \in S$, and thus $c_h^{n+1}(s,r) \sim_h (c_{hw_1}^n(s,r), ..., c_{hw_{|S|}}^n(s,r)) = (c_{hw_1}^{n-1}(s,r), ..., c_{hw_{|S|}}^{n-1}(s,r)) \sim_h c_h^n(s,r)$. □

Claim 2. *Let $|h|$ be the length of the history $h \in \mathcal{H}$. Then for every $n \geq 1$,*

$$c_h^n(s,r) = \mu(S_{n+|h|}S_{n+|h|+1} = sr|S_1...S_{|h|} = h).$$

In particular, $c_\emptyset^n(s,r) = \mu(S_nS_{n+1} = sr)$.

Proof of Claim 2. We prove the statement by induction on n, starting with the case base, $n = 1$. Indeed, $c_s(\mathbb{1}_r)\mathbb{1}_s = p_s(r)\mathbb{1}_r$. Thus, $c_\emptyset^1(s,r) = c_\emptyset(c_s(\mathbb{1}_r)\mathbb{1}_s) = p_\emptyset(s)p_s(r) = \mu(sr) = \mu(S_1S_2 = sr)$. The base case holds for every history following identical arguments.

Now, assume that the statement holds for n, and we prove it for $n+1$: $c_\emptyset^{n+1}(s,r) \sim_\emptyset (c_{w_1}^n(s,r), ..., c_{w_{|S|}}^n(s,r)) = \big(\mu(S_{n+1}S_{n+2} = sr|S_1 = w_1), ..., \mu(S_{n+1}S_{n+2} = sr|S_1 = w_{|S|})\big)$, where the last equality follows from the induction hypothesis. Thus, $c_\emptyset^{n+1}(s,r) = \sum_{w\in S} p_\emptyset(w)\mu(S_{n+1}S_{n+2} = sr|S_1 = w) = \mu(S_{n+1}S_{n+2} = sr)$. Again, the proof for non-empty histories follows identical arguments.

□

The following is an immediate corollary of the two claims above.

Corollary 1. *μ is 2-stationary.*

Claim 3. *$\mu(sr) = \mu(rs)$ for every $s, r \in S$.*

Proof of Claim 3. From Proposition 1 we know the posteriors converge with μ-probability 1. That is, in the limit the process determining the state realization is i.i.d. Thus, for $\epsilon > 0$ there exists $T > 0$ such that for every μ-positive probability $h \in S^T$, $\mu(S_T S_{T+1} = sr|h) = p_h(s)p_{hs}(r)$ is close up to ϵ to $\mu(S_T S_{T+1} = rs|h) = p_h(r)p_{hr}(s)$, for every $s, r \in S$. Since this is true for every history $h \in S^T$, then $\mu(S_T S_{T+1} = sr)$ is close up to ϵ to $\mu(S_T S_{T+1} = rs)$. From Corollary 1 above we know then that $\mu(sr) = \mu(S_1 S_2 = sr)$ and $\mu(rs) = \mu(S_1 S_2 = rs)$ are also close up to ϵ, but since ϵ is arbitrarily small, we have that $\mu(sr) = \mu(rs)$.

□

This completes the proof of the theorem.