THE DYNAMICS OF PREFERENCES, PREDICTIVE PROBABILITIES, AND LEARNING

EHUD LEHRER† AND ROEE TEPER‡

August 16, 2017

Abstract. We take a decision theoretic approach to predictive inference as in Blackwell [4, 5] and de Finetti [8, 9]. We construct a simple dynamic setup incorporating inherent uncertainty, where at any given time period the decision maker updates her posterior regarding the uncertainty related to the subsequent period. These predictive posteriors reflect the decision maker's preferences, period by period. We study the dynamics of the agents' posteriors and preferences, and show that a consistency axiom, reminiscent of the classic dynamic consistency, is the behavioral foundation for Bayesian updating according to noisy signals. We then focus on showing how such Bayesian updating is behaviorally distinct from exchangeability. It is pointed that even though beliefs of such Bayesian updating may not follow an exchangeable process, it seems as if they do from some point onwards.

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

JEL Classification: D81, D83

The authors would like to thank Shiri Alon-Eron, 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. Some of the results previously appeared in a manuscript circulated under the title "Who is a Bayesian?".

†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 in many economic models. In dynamic environments with uncertainty, as time goes by agents gather information, allowing them to update their perception of that uncertainty and infer regarding future events. When an agent is updating her belief in a Bayesian fashion, information might become increasingly useless as time passes. This phenomenon is referred to as *learning*: the agent learns as much as the dynamics allows her about the underlying uncertainty.[1] This paper takes a decision theoretic approach to Bayesian updating as in Blackwell [4, 5] and predictive inference as in de Finetti [8, 9]. We introduce a novel dynamic setup with inherent uncertainty and study the agent's preferences over uncertain alternatives, period by period.

In this set up we propose a behavioral property of preferences (an axiom) that is natural in this environment. On one hand it is reminiscent of the classical sure thing principle/dynamic consistency. On the other had, our decision environment is different than the classical one, and the axiom does not have the usual commitment interpretation. The axiom is termed *local consistency*. It turns out that this axiom is the behavioral foundation of the fundamental model introduced by Blackwell [4, 5], which in turn is equivalent to a martingale. In particular, over time learning occurs. This leads to the natural question of how such preferences differ from those who follow exchangeable processes (de Finetti [8]). The distinct behavioral features of exchangeable processes (relative to a general martingale) are shown. In addition, it is pointed that while these models are different and can be separated by axioms, the outcome of learning is similar.

1.1. **Bayesian updating: the local approach.** Consider a dynamic setup where at every period the decision maker (henceforth, DM) faces uncertainty regarding the upcoming state, and specifies her preferences over uncertain alternatives (Anscombe-Aumman [2] acts) for *the following period only*. Then, a state is realized and the DM makes up her mind regarding the next period. We assume that in every period 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 the posteriors and to find

[1] See Kalai and Lehrer [21], and Jackson, Kalai and Smorodinski [20] and references within for an extended study of learnability.

conditions under which different forms of Bayesian inference regarding future events take place. To gain a basic understanding we begin with a two-period model. That is, there is a single instance where the DM can accumulate new information. In this model the DM holds a prior belief over the state space S (representing the DM's preferences over Anscombe-Aumman acts for that period). Any realization of a state, say $t \in S$, at the first period, generates a posterior q_t over S (which typically differs from p). The posterior q_t represents the DM's preferences over Anscombe-Aumman acts for the second period.

As an example for how the posteriors are obtained from the prior, consider the classic model of Bayesian updating from a noisy signal that stochastically depends on the true state (as in Blackwell [4, 5]). For such a model in our environment the set S has two potential interpretations. First, it can be interpreted by the DM, and by a modeler as well, as the state space, which is the underlying payoff-relevant factor the DM is trying to asses. Second, for the DM, S can be interpreted as the set of signals she may observe. To explain this interpretation, suppose for instance, that there are only two states, low and high. Moreover, suppose the DM believes that in both states, she might observe 'high'. In this case, the signal 'high' tells her that the true state is high with some probability and low with the remaining probability. She may believe that in the two different states the signals are generated with different probabilities, and update her belief regarding the state that generated the signal accordingly. Having this in mind, the modeler interprets q_t as the DM's belief over the state space after observing t.

In our decision theoretic environment we introduce a condition on preferences referred to as *local consistency*. It states that if, at the second stage, preferences agree on the ranking of two alternatives, regardless of the information observed, then it must be that the two alternatives were ranked in the same way in the first stage as well. Stated differently, if the DM prefers to invest in firm A rather than in firm B, regardless of the information she observes today, then she should have preferred A over B also prior to observing any signal. Local consistency is reminiscent of the classical sure thing principle/dynamic consistency condition. However, our decision theoretic environment is different than the one in which these axioms are typically employed, and in particular the commitment interpretation is not suited for the current study.

We show that indeed local consistency has new implications. It is clearly necessary for Bayesian updating according to a noisy signal. Theorem 1 states that it is also sufficient. If local consistency is satisfied, one can elicit the probabilities according to which the DM expects to observe each signal, conditional on the state. In particular, when the DM updates her prior according to these probabilities, the posteriors derived from her preferences are obtained. We refer to systems of preferences, or beliefs, satisfying local consistency as *Blackwellian*.[2]

Local consistency has an additional representation, one that is useful when studying the infinite horizon model. It is well known that posteriors generated from noisy signals are a mean preserving spread of the prior. Theorem 1 shows that these two conditions are actually equivalent. That is, if the posteriors are a mean preserving spread of the prior, then it is possible to find noisy signals that would generate them from the prior through Bayesian updating.

1.2. **The global and local approaches.** Before introducing the infinite horizon model and the learning results, we point to the relation and distinction between the current approach and the classical one. In this paper we study updating in a local sense, similar to predictive inference. Here, the posteriors (and their dynamics) represent the DM's preferences (and their evolution) over outcomes determined *only* by the one-period-ahead signal realization. The posteriors naturally induce a *probability* over all (long and short) histories of signals. To clarify this statement, let p be the prior and q_t be the posterior following the history t. The probability μ over the two-period process is defined as $\mu(t, t') = p(t)q_t(t')$ (for every $t, t' \in S$). This probability is not directly revealed through the DM's preferences for bets over two-period state realizations; it is merely an auxiliary mathematical entity on which we can rely to make succinct statements regarding primitives. Our local approach stands in contrast with the global one, usually taken in the learning literature[3] and the updating literature (following Savage's [28] sure thing principle) discussing dynamic consistency.[4] In the global approach to dynamic consistency, the probability μ explicitly represents the primitives. The local approach introduced here gives rise to a different consistency axiom than the one prevalent in the existing literature. This axiom is one of the main issues of this study.

[2]This also assists in refraining from using the term "Bayesian," which may have a more general connotation than using Bayes rule upon observing a noisy signal.

[3]E.g., Epstein and Seo [13], Klibanoff, Mukerji and Seo [24] and others.

[4]See Epstein and Schneider [11], Sinischalchi [30] and Hanany and Klibanoff [18, 19].

1.3. **Martingales and exchangeable processes.** We extend the analysis from the two-period model to an infinite horizon one. From Theorem 1, the DM satisfies local consistency (following every history of signals) if and only if the process of posteriors forms a martingale. That is, when studying the dynamics of a Blackwellian system, one may alternatively study any martingale process and need not worry that some of the assumptions have been relaxed. This is useful since martingales are well understood processes with substantiated theory. For example, we can apply Doob's martingale convergence theorem and obtain that local consistency induces learning: the effect of additional information becomes increasingly negligible as the history observed by the DM becomes longer. Moreover, the posteriors converge.

Local consistency on its own does not impose any further restrictions on how the DM updates her beliefs, and may not satisfy many properties that are typically assumed in applications. For example, the signal structure following every history need not be the same as others, posteriors may not depend on the frequency of realizations, etc. Exchangeability (de Finetti [8]) is a particular form of martingales 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, for instance, that the empirical frequency of states must converge. Furthermore, the beliefs regarding subsequent states become closer and closer to the empirical frequency.

Exchangeable processes are classic and are frequently used in theory and applications. It is well known that an exchangeable process forms a martingale. It make studying the relation between martingales and exchangeable processes all the more interesting. In particular, what is the distinction between (general) martingales and exchangeability, and what behavioral feature makes an updating process exchangeable?

We provide two results. First, it is shown that local consistency is not only equivalent to a martingale, but implies a specific form of learning. It turns out, that processes that follow local consistency are *almost exchnageable*: 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 onwards it may seem to an outside observer that the DM closely followed an exchangeable process, and not just conducted Bayesian updating.

Second, we provide a condition on preferences implying that a martingale is also exchangeable. It turns out that the key behavioral difference between martingales and exchangeability is related to a property mentioned 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 even in the presence of local consistency the two requirements are not equivalent. Frequency dependence does not imply exchangeability. Example 2 shows that the fact that the posteriors following histories that have the same frequencies must coincide does not guarantee exchangeability. In the case of two states matters are relatively simple. Theorem 2 shows that in this case, frequency dependence, along with local consistency, axiomatize exchangeability. When more than two states are involved, frequency dependence is not strong enough: an additional consistency assumption (the formal discussion of which is deferred to the relevant section) is needed in order to characterize exchangeability. This is stated as Theorem 3.

The synthesis of our results implies the following. While local consistency induces different processes than exchangeable ones, the learning outcome is identical to that of an exchangeable process (Theorem 4). Now, whether a DM is interested in the precise evolution of beliefs right at the beginning, or only in the horizon, can be precisely determined by certain behavioral restrictions (Theorems 2 and 3).

1.4. **The structure of the paper.** The subsequent section introduces the two-period model and the main concepts. Section 3 extends the basic model to the infinite horizon and studies exchangeable processes. Section 4 introduces the notion of almost exchangeability and the relation of martingales to such processes. A discussion of the related literature and additional comments appear in Section 5. All proofs appear in the Appendix.

2. The Two-Period Model

In this section we introduce the decision theoretic environment and primitives. We start with a two-period model, which will later be extended to a dynamic model with infinite horizon.

2.1. **Acts.** We fix a finite *state space* S. An *act* is a real valued function defined over S taking values in $\mathbb{R}$. An act is interpreted as the utility the decision maker (DM) derives contingent on the realized state.[5] Denote the set of acts by $\mathcal{A}$. The classical theory of decision making under uncertainty[6] 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[7] $p \in \Delta(S)$ over the states S where an act $f \in \mathcal{A}$ is being evaluated according to its expected utility $\mathbb{E}_p(f) = \sum_{s \in S} p(s)f(s)$. We assume throughout that *every* preference relation discussed satisfies the Anscombe-Aumann assumptions and admits an SEU representation.

2.2. **Dynamics and preferences.** We are considering a dynamic environment in which the agent expresses her preferences over acts for the current period (only). Over time she may receive information regarding the underlying uncertainty and updates her beliefs. We start by analyzing a two-period model.

Formally, let $\succeq$ and $\{\succeq_s\}_{s \in S}$, be a *system* of preferences, each of which defined over $\mathcal{A}$. In the first period (*today*) the DM expresses her preferences, $\succeq$, over uncertain alternatives for the first period. Then, a state $s \in S$ is materialized, and in the second period (*tomorrow*) the DM is asked to express her preferences, $\succeq_s$, over acts for the

[5] 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.

[6] See, for example, Savage [28], Anscombe and Aumann [2], Schmeidler [29], Gilboa and Schmeidler [15], Karni [22], Bewely [3], Klibanoff, Marinacci, and Mukerji [23], Maccheroni, Marinacci and Rustichini [27], Chateauneuf and Faro [7], Lehrer and Teper [26], and many others.

[7] For a set A, denote by $\Delta(A)$ all Borel probability measures over A.

second period. We assume throughout that preferences admit an SEU representation. In particular, following every *history*, including the null one, we assume that the DM entertains some *posterior* over the state space S. We denote by p and p_s (for every $s \in S$) the system of posteriors representing $\succeq$ and $\succeq_s$, respectively.

The main question asked is what are the properties needed (beyond maximization of expected utility) to tie together the different history-dependent preferences. The goals are (a) to determine whether the beliefs the DM holds in the two periods regarding the underlying state are related in a natural sense, and (b) to figure out in what way the DM updates her belief about the future upon getting information between periods.

Bayesianism in the context of decision theory has been discussed broadly in the literature. The relation between an axiom (or a class of axioms) referred to as *dynamic consistency* and Bayesianism has been investigated (e.g., Hammond [16, 17], Chapter 7 in Tallon and Vergnaud [31] and the references within). Dynamic consistency is typically applied as follows. The underlying uncertainty is modeled globally": the DM is associated a preference over the entire duration of the model. More precisely, today the DM states her preferences over alternatives that are contingent acts; each such alternative is an act for today, and conditional on every possible piece of information between periods, an act for tomorrow. Then, in between periods a state is materialized, and the DM provides her preferences over acts for tomorrow.

Now consider the DM's ranking of two contingent acts, and assume they are identical accept for when a specific s is materialized in the first period. Dynamic consistency requires that her preferences today between such contingent acts, and her preferences tomorrow over the future acts, contingent on s materializing, have to be the same. That is, the DM's preferences today for alternatives tomorrow is consistent with her preferences tomorrow. With other classic axioms, dynamic consistency implies that the DM entertains an ex-ante belief μ over $S \times S$ according to which she assesses contingent acts today, and when s is materialized, she assesses alternatives in a Bayes consistent manner according to $\mu(\cdot|s)$.

In the current model, in contrast, we take a local approach. We assume that, given every history, the DM has preferences over choice objects defined over S, the stage-state space. In particular, in each period the DM states her preferences over one-period outcomes depending on the resolution of uncertainty in that period only. In our model the DM has merely a local view. This is the main conceptual difference between the current setup and those in the existing literature. True, a probability μ along with its conditionals can be defined over $S \times S$ by $\mu(s,t) = p(s)p_s(t)$. However, this equality is merely a mathematical entity which does not capture the DM's preferences and perception. Nevertheless, we are going to introduce this kind of extension in later sections in order to ease notation. It is important to note at this point that the construction of the probability μ places no further assumptions on preferences beyond expected utility, and cannot guarantee that the posteriors themselves evolve according to an updating with respect to some noisy information structure.

2.3. **Local consistency.** We introduce the following property, which postulates that if regardless of the state realized today the DM prefers f over g tomorrow, then today she should prefer f over g as well. Formally:

Local Consistency. For every $f, g \in \mathcal{A}$, if $f \succeq_s g$ for every $s \in S$, then $f \succeq g$.

Local consistency is reminiscent of dynamic consistency, but since the approach we take is local, it does not imply any commitment on the part of DM as in the global approach. In addition, dynamic consistency puts no restrictions on beliefs beyond subjective expected utility, while *local consistency* clearly caries additional behavioral content. It points to the idea that the underlying uncertainty today and that of tomorrow are related, and that the DM should expect the information received in between periods to at least partially convey this relation. Indeed, if the posteriors representing $(\succeq_s)_{s\in S}$ are a mean preserving spread of the belief representing $\succeq$, then the condition holds.[8] In environments for which the uncertainty today and tomorrow are not related, this axiom is less appealing. A Markov process is a good example. If state s is realized, then the transition probabilities for tomorrow conditional on s do not depend on what is the law under which state s was chosen today. In the next sections we see exactly how these intuitions are formalized.

Note, the axiom can also be interpreted as a no "Dutch Book" condition; if regardless of the information about the realization of uncertainty the DM will choose tomorrow f over g, then she should not be swayed to choose g over f today, and tomorrow swap back to f for sure.

2.4. **Blackwell's comparison.** Blackwell's comparison of experiments [4, 5] is classical in the theory of Bayesian updating. A DM has a prior probability regarding the

[8] $(q_s)_{s\in S}$ is a mean preserving spread of p if there is some distribution $\alpha \in \Delta(S)$ such that $p = \sum_{s\in S} \alpha_s q_s$.

state of nature. Before she takes a decision, she receives a noisy signal that depends on the true state. Formally, a stochastic matrix reflecting what is the likelihood of a signal conditional on the parameter is given. Once a signal is observed, the DM updates the prior according to the matrix and Bayes rule to obtain a posterior, and takes a decision.

We integrate this idea in our framework. In our model the set of signals coincides with the state space S: any $s \in S$ is referred to not only as a state but also as a signal. For example, there are two possible states of the economy, low and high. Then a signal, whether the true state is high or low, is observed (say, through the return from an investment). Since the signal stochastically depends on the true state, it partially reveals the identity of true state of the economy.[9]

Formally, let $p \in \Delta(S)$ be some prior over the state space and π a stochastic map—$\pi(t|s)$ is the probability to observe t conditional on s. Assume that signal $t \in S$ is observed. By Bayes updating the posterior obtained is, $q_t^{p,\pi} \in \Delta(S)$ defined by

$$q_t^{p,\pi}(s) = \frac{p(s)\pi(t|s)}{\sum_{s'} p(s')\pi(t|s')}.$$

The primitive in our decision theoretic setup is, $p, (q_t)_{t \in S}$, a system of beliefs, or equivalently a system of SEU preferences, each of which is represented by a belief over the state space S. We ask when is it that the system is such that there exists a stochastic map π, where $q_t = q_t^{p,\pi}$ for every $t \in S$. In words, under what conditions the system can be explained by Bayes updating of p with respect to some stochastic map π?

Definition 1. *Let $p, (q_t)_{t \in S}$ be a system. We say that the system is* Blackwellian *if there is a stochastic map π such that $q_t = q_t^{p,\pi}$ for every $t \in S$.*

2.5. **Preferences and Blackwellian systems.** Consider a system of preferences $\succeq$ and $(\succeq_t)_{t \in S}$. We interpret it as follows: the DM has to convey her preferences over uncertain alternatives before ($\succeq$) and after ($\succeq_t$) the realization of a state $t \in S$. However, we do not make an assumption that the DM herself interprets it in the same way as us, the modelers. She may believe that what is being observed between periods, namely $t \in S$, is merely a signal of what the true state is, and not the state itself. This also explains why in the model of Blackwellian systems the state space and the signal

[9]The fact that we identify states with signals is an artifact of our decision theoretic environment. We elaborate on this point prior to the main preference representation result of this section.

space are identical; in our set up the modeler and the DM observe the same piece of information, but they may interpret it differently.

With this in mind, we are ready to present the main representation result of this section.

Theorem 1. *Let* $\succeq$ *and* $(\succeq_t)_{t\in S}$ *be a system with* $p, (q_t)_{t\in S}$ *representing posteriors. Then the following are equivalent:*

(1) Local consistency *is satisfied;*
(2) *The system is Blackwellian; and*
(3) *There is an* $\alpha \in \Delta(S)$ *such that* $p = \sum_{t\in S} \alpha_t q_t$.

Furthermore, if π *is a subjective stochastic map such that* $q_t = q_t^{p,\pi}$ *for every* $t \in S$, *then* α *can be obtained by* $\alpha_t = \sum_{s\in S} p(s)\pi(t|s)$ *for every* $t \in S$. *Alternatively, if* $\alpha \in \Delta(S)$ *is such that* $p = \sum_{t\in S} \alpha_t q_t$, *then for* π *defined by* $\pi(s|t) = \alpha_t \frac{q_t(s)}{p(s)}$ *for every* $t, s \in S$, *we obtain* $q_t = q_t^{p,\pi}$.

The theorem identifies *local consistency* as a necessary and sufficient condition for a system to be Blackwellian. It means that this condition is met if, and only if, there is a subjective stochastic map according to which the DM updates p to obtain q_t whenever t is observed, for every $t \in S$.

It is well known that a system being Blackwellian implies that the q_t's are a mean preserving spread of p. The theorem states that the opposite implication is also true. If we consider a set of beliefs $(q_t)_{t\in S}$ that is a mean preserving spread of p, then it must be that the system is Blackwellian. This part of Theorem 1 will be helpful in the subsequent section when we study the infinite horizon model.

3. The Infinite Horizon Model

We now turn to analyze the infinite horizon model. We further study Blackwellian systems, and focus on implications to learning. Consider the 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, like before, 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.

We expand the model presented in the previous section. 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 that every history-dependent preferences $\succeq_h$ admit an SEU representation with respect to the *posterior* $p_h \in \Delta(S)$.

3.1. **Infinite histories and the generated probability.** Let $\Omega = S^{\mathbb{N}}$ be the set of 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 $\omega \in h$.

Endow the set of infinite histories Ω 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 will be convenient to think about the evolution of the posteriors $\{p_h\}_{h\in\mathcal{H}}$ by means of the probability μ they generate.

3.2. **Local consistency and martingales.** Extending our theory from the previous section to the infinite horizon model, we first consider preferences that are locally consistent regardless of the history. That is, for every history $h \in \mathcal{H}$ and $f, g \in \mathcal{A}$, if $f \succeq_{hs} g$ for every $s \in S$, then $f \succeq_h g$. We abuse terminology and refer to the same term used in the two-period model. The interpretation remains.

Theorem 1 implies that *local consistency* is satisfied if and only if for every history $h \in \mathcal{H}$ the system associated with the histories h and $(hs)_{s\in S}$ is Blackwellian. That is, following every history, the system follows updating according to Bayes rule with respect to some subjective stochastic map. Note that these matrices could differ as a function of the history. We can also deduce that this is equivalent to the system, following every history, being a mean preserving spread. That is, for every non-null history h there is an $\alpha^h \in \Delta(S)$ such that $p_h = \sum_{s\in S} \alpha_s^h p_{hs}$. Now, let $\alpha \in \Delta(\Omega)$ be the distribution over Ω with α-probability of state $s \in S$ conditional on history $h \in \mathcal{H}$

equals α_s^h. Then, for every history h we obtain

$$\mathbb{E}_{\alpha}\left(p_{hs}|h\right) = p_h. \tag{1}$$

We see that Theorem 1 has broad implications to the infinite horizon model. It shows that Blackwellian systems are exactly the well studied class of models in the learning literature, namely martingales. In particular, the system of posteriors form a martingale with respect to the distribution α. Since our interest is in Blackwellian systems, this implies that we can focus on martingales, benefitting from the familiar tools in that literature, without being concerned that it is a too general model. We will alternate between using the equivalent terms Blackwellian systems and martingales.

Having this in mind, we turn to investigating a particular class of martingales well studied and applied in the statistics and economics literatures, namely exchangeable processes. The main question is what are the behavioral properties, if any, distinct such processes from general Blackwellian processes.

3.3. **Exchangeability.** Consider a DM who believes that a distribution (or, a parameter) p over the state space S is chosen randomly according to some distribution $\theta \in \Delta(\Delta(S))$. Assume that states are selected according to the chosen p in an i.i.d. manner. Also assume that the DM does not know what is 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 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.[10] The following definition is due to de Finetti [8]:

Definition 2. *μ is* exchangeable *if whenever two histories $h, h' \in \mathcal{H}$ share the same frequency, $\mu(h) = \mu(h')$.*

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

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.*"

So, how exactly exchangeable processes differ from martingales? For example, while in exchangeable processes the DM obtains signals that are i.i.d. conditional on the state-generating parameter, a general Blackwellian considers an abstract signal structure (that may depend on the history of realizations). But perhaps there is a corresponding representation in terms of the de Finetti representation?

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 system of posteriors forms a martingale and thus local consistency *is satisfied. The process itself however is not exchangeable since the probability of HTT is different than that of THT.*

Exchangeable processes are martingales and thus satisfy our *local consistency* axiom. The example above shows that not every martingale is exchangeable. In the subsequent sections we find conditions in terms of the primitives that identify which martingale is indeed an exchangeable process.

3.4. **Self confirming processes.** As discussed in Section 2, when the DM observes the state realization, it is as if she believes it is merely a signal regarding what the state was. In exchangeable processes however, the information is the sate itself, and the DM believes this as well.

So consider a system $p_h, (p_{hs})_{s\in S}$ following some history $h \in \mathcal{H}$. If *local consistency* is satisfied, then this system is Blackwellian and there is an $\alpha \in \Delta(S)$ such that $p_h =$

$\sum_{s\in S}\alpha_s p_{hs}$. If $\alpha_s \neq p(s)$ for some $s \in S$, then the DM assigns different probabilities to event that state s occurring and to the event of observing s. Thus, a necessary condition for a martingale being exchangeable, is that $\alpha_s = p(s)$. In this case we obtain that $\sum_{s\in S} p_h(s)p_{hs}(t) = p(t)$ for every $t \in S$, and writing it as in Eq. (1), we obtain,

$$\mathbb{E}_\mu\left(p_{hs}|h\right) = p_h.$$

We will refer to systems satisfying this condition as *self confirming.*[11] Note, that this is not a sufficient condition for exchangeability; the martingale in Example 1 satisfies this property, but is not exchangeable.

We now turn to formalizing this condition in terms of preferences. 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 the DM perceives f, looking one period ahead into the future. The following axiom postulates that, given a 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.

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. 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, a history is null if its μ-probability is 0.

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

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

[11] There is a minor issue of μ-probability 0 histories, but the axiom and formal result we will present take this into consideration.

Strong local consistency is a formulation of the idea that if the state/signal is not going to be observed, the DM's assessment of acts should not change. In other words, the underlying uncertainty in the future should not be assessed differently than in the present, if the information regarding the underlying uncertainty is a counterfactual. This axiom is appealing, just like *local consistency*, when the underlying uncertainty tomorrow is related to today. However, it is conceptually stronger in the sense that absent any information, preferences should not change. Note, the axiom can be reformulated as in *local consistency* and postulate that for every non-null history h and acts f, g,

$$f \succeq_h g \text{ if and only if } \hat{c}_h(f) \succeq_h \hat{c}_h(g).$$

From the latter formulation we can see that it is indeed stronger version of *local consistency.* This is true even without the assumption that preferences admit an SEU representation. All is needed is that every act has a certainty equivalent and that preferences are monotonic.[12] Indeed, if for some history h and acts f, g we have $f \succeq_{hs} g$ for every $s \in S$, then $c_{hs}(f) \succeq_{hs} c_{hs}(g)$ for every $s \in S$. If preferences are monotonic, then it is implied that $\hat{c}_h(f) \succeq_h \hat{c}_h(g)$.

Proposition 1. *Let $(\succeq_h)_{h \in \mathcal{H}}$ be a system with representing posteriors $(p_h)_{h \in \mathcal{H}}$. Then the following are equivalent:*

(1) Strong local consistency *is satisfied; and*
(2) *The system is Blackwellian and self confirming. In particular,*

$$\mathbb{E}_\mu \left(p_{hs} | h \right) = p_h \quad \mu\text{-a.s.}$$

The proposition has implications beyond bringing us one step closer to understanding the behavioral underpinning of excangeability in our setup, and how it differs from general Blackwellian systems. From the martingale convergence theorem (Doob [10]) we know that the posteriors converge with μ-probability 1. As discussed above, 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. The proposition (and subsequent results) can be interpreted as follows: if *strong local consistency* is satisfied then the modeler knows the DM is going to learn, where learning is accomplished with respect to the natural measure μ, which is consistent with the DM's beliefs. In a similar manner, the modeler can appeal to the DM's reasoning. If she

[12] A preference $\succeq$ is monotonic if whenever $f(s) \succeq g(s)$ for every $s \in S$, then $f \succeq g$.

satisfies *strong local consistency*, then she is guaranteed to learn with respect to the natural measure induced by her beliefs. Subsequent results may be interpreted in a similar fashion.

3.5. **Frequency-determined preferences and Exchangeability.** We now provide a condition that guarantees not only that the DM is a self confirming Blackwellian, but also that her posteriors follow a learning pattern akin to an exchangeable process.

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

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 2. *Let $S = \{H, T\}$. The first two periods of the process are as in Example 1: 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* strong 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 *strong 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,* strong local consistency *and* frequency dependence *are satisfied if and only if μ is exchangeable.*

3.6. **Spaces with more than two states.** 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 3 below shows that when S consists of more than two states, the probabilities of HT and TH are not necessarily equal even though *strong local consistency* is satisfied.

Example 3. *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* strong 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, $\nu((S_1...S_k) = h) = \nu((S_t S_{t+1}...S_{t+k-1}) = h)$ for any time t. When a process is k-stationary for every k we say that it is *stationary.* While *strong 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 3 continued. *The main point of this example is that $\mu(S_1 S_2 = HT) = 0$, while $\mu(S_2 S_3 = HT) = \frac{1}{4} \cdot \frac{1}{3} \cdot \frac{2}{3} + \frac{1}{4} \cdot \frac{2}{3} \cdot \frac{1}{3} > 0$. In particular, $\mu(S_1 S_2 = HT) \neq \mu(S_2 S_3 = HT)$, making the process non-stationary.*

From Example 3 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 *strong 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 3.4, $c_{hs}(\mathbb{1}_r)$ 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 *strong 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 occurred), and a bet on s tomorrow followed by a bet on r in the following period (provided that s has indeed occurred).

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

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

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

Remark 1. *The preferences in Example 3 satisfy* strong 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, strong local consistency *implies 2-stationarity (and exchangeability from Theorem 3), yet this question remains open at the current point.*

4. Martingales and Almost Exchangeable Processes

In Section 3.4 we discussed how since a system of posteriors satisfying *strong local consistency* is a self confirming martingale, it must be from the martingale convergence theorem that the posteriors converge with probability 1. In the current section we show that in the long run any such martingale is as close as we wish to an exchangeable process, even if it does not satisfy *frequency dependence* (and *two-tier local consistency*). Initially the updating need not be identical to an exchangeable process. Nevertheless, there exists an exchangeable process such that from some point onwards, the set of histories, for which the posteriors are close to those associated with the exchangeable process, is of probability close to 1.

Consider now a processes of posteriors $\{p_h\}_{h\in\mathcal{H}}$ converging with μ-probability 1, and let $C \subseteq \Omega$ be the set of (infinite) histories of μ-probability 1 for which the posteriors converge. That is,[13] for every $\omega \in C$, p_{ω^T} converges (as $T \to \infty$) to a limit denoted as p_ω. Thus, $\{p_\omega\}_{\omega\in C}$ and μ induce a probability distribution over $\Delta(S)$, which we denote as θ. Note that martingales fall in this realm of processes. Also note that if μ is exchangeable, then θ completely characterizes the evolution of the posteriors, and not just the limits as does a martingale.

For a given probability (generated by a family of history dependent posteriors), let S_T be the random variable standing for the realized state at time T.

Definition 3. *Let μ and μ' be two probabilities, each generated by a processes of posteriors converging with μ and μ' probability* 1*, respectively. Let θ and θ' be the respective induced distributions over $\Delta(S)$. We say that μ and μ' are ϵ-*close *if*
(i) $|\theta(A) - \theta'(A)| < \epsilon$ for every measurable $A \subseteq \Delta(S)$; and
(ii) $|\mu(S_\tau = s) - \mu'(S_\tau = s)| < \epsilon$ for every $s \in S$ and every time τ.

The definition states that two probabilities are close up to ϵ, if (i) the distributions over the limits of the two martingales are close up to ϵ, and (ii) in every period, the distributions over the realized state, induced by the two processes, are also close up to ϵ.

[13]For $\omega = (\omega_1, \omega_2, \dots) \in \Omega$ and T, denote by $\omega^T = (\omega_1, \dots, \omega_T)$ the prefix of ω whose length is T.

Definition 4. *The probability μ is* almost exchangeable *if for every $\epsilon > 0$ there exists an exchangeable probability ζ_ϵ such that*

(a) *there exists $T > 0$ satisfying*

$$|\mu(hs|h) - \zeta_\epsilon(hs|h)| < \epsilon \text{ for every } s \in S$$

for a set of histories $h \in S^T$ whose μ-probability is at least $1 - \epsilon$; and

(b) *μ and ζ_ϵ are ϵ-close.*

Theorem 4. *If the posteriors form a self confirming martingale, then the probability is almost exchangeable.*

The following example illustrates Theorem 4 and conveys some of the intuition behind this result.

Example 1 continued. *While it may seem non-intuitive at a first glance, the process in Example 1 is almost exchangeable. To show why it could look non-intuitive, consider two (very long) histories in which T appears only once, and they differ from one another by whether T appears at the first period or the last: $h = (T, H, ..., H)$ and $h' = (H, ..., H, T)$. Obviously, an exchangeable process, say ζ, induces the same posterior after both h and h'. That is, $\zeta(hH|h) = \zeta(h'H|h')$. If our statement (that the process μ in the example is almost exchangeable) is correct, then $\zeta(hH|h)$ must be very close to $\mu(hH|h) = \frac{1}{3}$. At the same time $\zeta(h'H|h')$ must be very close to $\mu(h'H|h') = \frac{2}{3}$. This is impossible. While the reasoning applied to h and h' is correct, it is misleading. The reason is that histories (like h and h') indicating that our statement is counter intuitive cannot both have non-negligible probability. In the long-run, a long stretch of H's induces a high posterior probability on (the next outcome being) H. Thus, if h has a non-negligible probability, then the probability of h' must be negligible. In what follows we show that the process in Example 1 is indeed almost exchangeable.*

Consider the following exchangeable process ζ: with probability $\frac{1}{2}$ it chooses a biased coin assigning probability $\frac{2}{3}$ to H and with probability $\frac{1}{2}$ it chooses a biased coin assigning probability $\frac{1}{3}$ to H. As for μ, if H is the outcome of the first toss, the law of large numbers implies that the frequency of H will converge to $\frac{2}{3}$ with μ probability 1. But this is also true if in producing the distribution of ζ, the coin chosen is the $\frac{2}{3}$-biased. The posteriors related to ζ converge to $\frac{2}{3}$. Thus, while not necessarily true early on, as time goes by the posteriors associated to μ and those associated to ζ are getting closer to each other (on a collection of histories of μ-probability that is getting close to 1). A *similar argument holds in the case where the outcome of the first toss determining* μ *is* T*.*

Note that in this example all the exchangeable processes ζ_ϵ in the definition of almost exchangeable coincide with ζ. The next example shows that this is not always possible.

Example 4. *Consider the process* μ *in which we throw a fair coin in the first period; if the outcome is* H *we forever throw a coin that yields* T *for sure, and if the outcome in the first period is* T *we forever throw a coin that yields* H *for sure. This process is a martingale. The distribution it generates over* $\Delta(\{H, T\})$ *is* 0.5 *on the parameter that yields* H *for sure, and* 0.5 *on the parameter that yields* T *for sure. However* μ *does not approximate* ζ *according to Definition 4. While item* 2 *of the definition holds, item* 1 *fails since it is impossible under* ζ *that* T *would follow an* H *(and vice versa).*

The way to rectify this, and this construction is explicit in the proof of Theorem 4, is to ϵ *perturb the parameters of* ζ *and let* ζ_ϵ *as the process defined by the parameters of* ζ *when they are* ϵ*-mixed with a uniform distribution over* $\Delta(\{H, T\})$*.*

5. The Literature and Additional Comments

5.1. **Literature.** Several papers have studied different aspects of Bayesianism and exchangeability in the context of decision theory (e.g., Epstein and Seo [13], Klibanoff, Mukerji and Seo [24] 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.[14] In such a setup, axiomatizing exchangeability in the context of expected utility is a straightforward application of de Finetti's [8] characterization and is obtained by assuming symmetry of preferences for (finite) permutations of the experiments' outcomes.

Epstein and Schneider [11], Sinischalchi [30] and Hanany and Klibanoff [18, 19] 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 [12] without an axiomatic foundation.

[14]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 [14] consider processes of posteriors, where the primary concern is to identify when such processes are exchangeable. This is as opposed to the current study, focusing on general Bayesianism and its relation to exchangeability. Our results do not imply, nor are implied by, the results in [14]. One of their properties is similar to our *frequency dependence* when translated to posteriors. They introduce a second condition which does not characterize martingales (or Bayesianism) and is a strengthening of a two-symmetry condition we discuss in the proofs of Theorems 2 and 3.

5.2. **Symmetry in the current framework.** It is possible to use the 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).[15] This seems a much more complicated task.

Lastly, our approach separates (self confirming) Blacwellian systems and exchangeable processes. This woud not have been possible if we axiomatized exchangeable processes directly through a symmetry axiom.

[15]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*.

5.3. **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 2, however, the inverse direction is typically incorrect.

The consistency axioms (i.e., *strong local consistency* and *two-tier local consistency*) imply 1 and 2-stationarity. But from Proposition 1, *strong local 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*, 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.

5.4. **Lexicographic probability systems.** By assuming that *strong local consistency* even for non-null histories, it is possible to reformulate the results and to incorporate learning 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 *strong local consistency* is satisfied for that history as well, the probability generated by the posteriors associated with the continuation histories (of history h) is a self confirming martingale.[16] Similar modifications can be made to the other results.

[16]This construction is similar to the one behind *Lexicographic Probability Systems* presented in Brandenburger, Friedenberg and Keisler [6] (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.

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] Blackwell, D. (1951) Comparison of Experiments, *Proceedings Second Berkeley Symposium on Mathematics, Statistics and Probability, University of California Press*, 93102.

[5] Blackwell, D. (1953) Equivalent Comparison of Experiments, *Annals of Mathematics and Statistics*, **24**, 265272.

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

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

[8] 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.

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

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

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

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

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

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

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

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

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

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

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

[20] Jackson, M. O., E. Kalai, and R. Smorodinsky (1999) Bayesian Representation of Stochastic Processes under Learning: de Finetti Revisited, *Econometrica*, **67**, 875–893.

[21] Kalai, E., and E. Lehrer (1994) Rational Learning Leads to Nash Equilibrium, *Econometrica*, **61**, 1019–1045.

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

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

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

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

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

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

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

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

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

[31] 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 Theorem 1.** We start by showing that *local consistency* is equivalent to $(q_t)_{t\in S}$ being a mean preserving spread of p. *local consistency* is clearly necessary. To see that it is sufficient, assume that *local consistency* and that p is not in the convex hull of $(q_t)_{t\in S}$. Thus we can find a separating vector $f \in \mathbb{R}$ such that $\mathbb{E}_p(f) > 0 = E_p(0)$ while $\mathbb{E}_{q_t}(f) \geq 0 = E_{q_t}(0)$ for every $t \in S$. But the later contradicts *local consistency.*

We now turn to show that $(q_t)_{t\in S}$ being a mean preserving spread of p is equivalent to the system being Blackwellian. Again, $(q_t)_{t\in S}$ being a mean preserving spread of p is clearly necessary. We show that it is sufficient. Let $\alpha \in \Delta(S)$ such that $p = \sum_{t\in S} \alpha_t q_t$. Now, define π as follows: $\pi(t|s) = \alpha_t \frac{q_t(s)}{p(s)}$. If we show that π is indeed a stochastic map and that $q_t = q_t^{p,\pi}$, then the proof is complete. Indeed,

$$\sum_{t\in S} \pi(t|s) = \sum_{t\in S} \alpha_t \frac{q_t(s)}{p(s)} = \frac{1}{p(s)} \sum_{t\in S} \alpha_t q_t(s) = \frac{1}{p(s)} p(s) = 1,$$

for every $s \in S$. Lastly, for every $t, s \in S$ we obtain

$$q_t^{p,\pi}(s) = \frac{p(s)\pi(t|s)}{\sum_{s'} p(s')\pi(t|s')} = \frac{q_t(s)\alpha_t}{\alpha_t} = q_t(s).$$

A.2. **Proof of Proposition 1.** In order to see the sufficiency of *strong local consistency*, fix a history h and let $\mathbb{1}_E$ denote the indicator function for an event $E \subseteq S$. We know that *strong local consistency* 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 *strong local consistency* is immediate and is thus omitted.

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

$$p(1-q) = (1-p)r. \tag{2}$$

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 *strong local consistency*, which means that Eq. (2) 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,

$$\phi(\bar{h}) = \phi(\bar{h}'). \tag{3}$$

From the induction assumption, $\mu(\bar{h}) = \mu(\bar{h}')$. Also implied from Eq. (3) 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 *strong 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

$$\mu(\bar{h}s) = \mu(\hat{h}ts) \text{ and } \mu(\bar{h}'t) = \mu(\hat{h}'st). \tag{4}$$

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. (4) 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. (4) 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

$$c_h^{n+1}(s,r) \sim_h (c_{hw_1}^n(s,r), ..., c_{hw_{|S|}}^n(s,r)). \tag{5}$$

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 $c^n_{hw}(s,r) \sim_{hw} c^{n-1}_{hw}(s,r)$ for every $w \in S$, and thus $c^{n+1}_h(s,r) \sim_h (c^n_{hw_1}(s,r),...,c^n_{hw_{|S|}}(s,r)) = (c^{n-1}_{hw_1}(s,r),...,c^{n-1}_{hw_{|S|}}(s,r)) \sim_h c^n_h(s,r)$.

□

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

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

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

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, $c^1_\emptyset(s,r) = c_\emptyset(c_s(\mathbb{1}_r)\mathbb{1}_s) = p_\emptyset(s)p_s(r) = \mu(sr) = \mu(S_1S_2 = sr)$. The base case holds for every history following identical arguments.

Now, assume that the statement holds for n, and we prove it for $n+1$: $c^{n+1}_\emptyset(s,r) \sim_\emptyset (c^n_{w_1}(s,r),...,c^n_{w_{|S|}}(s,r)) = \big(\mu(S_{n+1}S_{n+2} = sr|S_1 = w_1),...,\mu(S_{n+1}S_{n+2} = sr|S_1 = w_{|S|})\big)$, where the last equality follows from the induction hypothesis. Thus, $c^{n+1}_\emptyset(s,r) = \sum_{w\in S} p_\emptyset(w)\mu(S_{n+1}S_{n+2} = sr|S_1 = w) = \mu(S_{n+1}S_{n+2} = sr)$. 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 $h \in S^T$, $\mu(S_TS_{T+1} = sr|h) = p_h(s)p_{hs}(r)$ is close up to ϵ to $\mu(S_TS_{T+1} = rs|h) = p_h(r)p_{hr}(s)$, for every $s,r \in S$. Since this is true for every history $h \in S^T$, then $\mu(S_TS_{T+1} = sr)$ is close up to ϵ to $\mu(S_TS_{T+1} = rs)$. From Corollary 1 above we know then that $\mu(sr) = \mu(S_1S_2 = sr)$ and $\mu(rs) = \mu(S_1S_2 = rs)$ 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.

A.5. **Proof of Theorem 4.** 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 *strong local consistency*, then by Proposition 1 the posteriors form a martingale. Thus, the posteriors converge μ-almost surely (see Doob [10]). 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 $t_1(\varepsilon)$ large enough so that[17] for every $T \geq t_1(\varepsilon)$

$$\mu\left(\omega \in C;\ ||p_\omega - p_{_{\omega^T}}|| < \varepsilon/3\right) > 1 - \varepsilon/3. \tag{6}$$

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 the set of ω's that do not share this property, which is measurable.)

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

$$\mu\left(\omega \in \Omega;\ \frac{|\sum_{t=t_1(\varepsilon)}^{t_2}[\mathbb{1}_{\omega_t = s} - p_{_{\omega^t}}(s)]|}{t_2 - t_1(\varepsilon) + 1} < \varepsilon/3 \text{ for every } s \in S \;\middle|\; \omega^{t_1(\varepsilon)}\right) > 1 - \varepsilon/3. \tag{7}$$

Recall that for $\omega \in C$, all posteriors, $p_{_{\omega^T}}$, for $T > t_1(\varepsilon)$, are close to p_ω up to $\varepsilon/3$. Denote by $\phi(h)$ the relative frequency of history h. Due to (6) we obtain,

$$\mu\left(\omega \in C;\ ||\phi(\omega^{t_1(\varepsilon),t_2}) - p_{_\omega}|| < 2\varepsilon/3\right) > 1 - 2\varepsilon/3. \tag{8}$$

If $t_2(\varepsilon)$ is large enough, then $\frac{t_1(\varepsilon)}{t_2(\varepsilon)} < \varepsilon/3$ and the first $t_1(\varepsilon)$ states in $\omega^{t_2(\varepsilon)}$ have a weight smaller than $\varepsilon/3$. Thus, (8) changes to

$$\mu\left(\omega \in C;\ ||\phi(\omega^{t_2}) - p_\omega|| < \varepsilon\right) > 1 - 2\varepsilon/3. \tag{9}$$

In words, the probability of $\omega \in C$ such that the empirical frequency of the states over the history ω^{t_2} (recall, $t_2 \geq t_2(\varepsilon)$) 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

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

histories ω^{t_2} and their posteriors are close to each other. Formally, with the help of (6), when $t_2 \geq t_2(\varepsilon) > 3t_1(\varepsilon)/\varepsilon$,

$$\mu\left(\omega \in C;\ ||\phi(\omega^{t_2}) - p_{_{\omega^{t_2}}}|| < 4\varepsilon/3 \text{ and } ||\phi(\omega^{t_2}) - p_\omega|| < \varepsilon\right) > 1 - 2\varepsilon/3. \tag{10}$$

So much for μ.

We turn now to the definition of an exchangeable process, ζ_ϵ. 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 $(\Delta(S), \theta)$. Now, let U be a uniform distribution over $\Delta(S)$, and for $\epsilon > 0$ let θ_ϵ be $\epsilon U + (1-\epsilon)\theta$. For every $p \in \Delta(S)$ let B_p be the i.i.d. process with stage-distribution p. ζ_ϵ is the process defined by the distribution over $\{B_p; p \in \Delta(S)\}$ induced by the distribution θ_ϵ over the p's (i.e., over $\Delta(S)$).

We start by showing that item (b) of Definition 4 holds. The first part is immediate by the definition of ζ_ϵ. Since the distributions over $\Delta(S)$, θ and θ_ϵ, induced respectively by μ and ζ_ϵ are ϵ-close, in the limit, the distribution over the realized state is close up to ϵ. Since both μ and ζ_ϵ are martingales, this is true for every time period t.

We move to show item (a) of Definition 4 holds. The definition of ζ_ϵ and the fact that it is exchangeable imply the following facts:
(i) For ζ_ϵ-almost every $\omega \in \Omega$, the limit of the empirical frequencies, $\lim_{t\to\infty} \phi(\omega^t)$, 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_ϵ-almost every $\omega \in \Omega$, the posteriors w.r.t. to ζ_ϵ, denoted $p^{\zeta_\epsilon}_{\omega^t}$, converge to q_ω as $t \to \infty$. That is, $\lim_{t\to\infty} p^{\zeta_\epsilon}_{\omega^t} = q_\omega$ with ζ_ϵ-probability 1.
(iii) The distributions of p_ω (induced by μ) and that of q_ω (induced by ζ_ϵ) (both over $\Delta(S)$) are close up to ε as appears in Definition 3 item (i).
(iv) For θ-almost every $p \in \Delta(S)$ with ζ_ϵ high probability, if the relative frequency $\phi(\omega^T)$ is close to p, then the posterior (the one induced by ζ_ϵ 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 $t > t_3(\varepsilon)$,

$$\zeta\left(\omega;\ ||\phi(\omega^t) - p|| < \varepsilon/2 \text{ implies } ||p^{\zeta_\epsilon}_{\omega^t} - p|| < \varepsilon\right) > 1 - \varepsilon/3. \tag{11}$$

In other words, 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 $t > t_3(\varepsilon)$,

$$||\phi(\omega^t) - p|| < \varepsilon/2 \text{ implies } ||p^{\zeta_\epsilon}_{\omega^t} - p_\omega|| < \varepsilon. \tag{12}$$

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

$$\mu\left(\omega;\ \text{there is } p \in \Delta(\varepsilon) \text{ such that } ||\phi(\omega^{t_4}) - p|| < \varepsilon/2\right) > 1 - \varepsilon/3. \tag{13}$$

From (10) and (13) we obtain that for every $t > \max\{t_2(\varepsilon), t_4(\varepsilon)\}$,

$$\mu\Bigg(\omega \in C;\ ||\phi(\omega^t) - p_{_{\omega^t}}|| < 4\varepsilon/3, ||\phi(\omega^t) - p_\omega|| < \varepsilon \\ \text{and there is } p \in \Delta(\varepsilon) \text{ such that } ||\phi(\omega^t) - p|| < \varepsilon/2\Bigg) > 1 - \varepsilon. \tag{14}$$

Combining with (12) we get that when $t > \max\{t_2(\varepsilon), t_3(\varepsilon), t_4(\varepsilon)\}$,

$$\mu\Bigg(\omega \in C;\ ||\phi(\omega^t) - p_{_{\omega^t}}|| < 4\varepsilon/3, ||\phi(\omega^t) - p_\omega|| < \varepsilon \\ \text{and } ||p^{\zeta_\epsilon}_{\omega^t} - p_\omega|| < \varepsilon\Bigg) > 1 - 4\varepsilon/3. \tag{15}$$

Thus (due to the triangle inequality),

$$\mu\left(\omega \in C;\ ||p^{\zeta_\epsilon}_{\omega^t} - p_{_{\omega^t}}|| < 10\varepsilon/3\right) > 1 - 4\varepsilon/3. \tag{16}$$

This completes the proof.