



Exploration and correlation [☆]

Evan Piermont ^{a,*}, Roe Teper ^b



^a Royal Holloway, University of London, Department of Economics, United Kingdom of Great Britain and Northern Ireland

^b University of Pittsburgh, Department of Economics, United States of America

ARTICLE INFO

Article history:

Received 9 January 2019

Available online 23 April 2019

JEL classification:

D81

D83

Keywords:

Bandit problems

Correlated arms

Strong exchangeability

ABSTRACT

We note that in environments such as exploration problems, in which agents have to choose a single action out of several in each period, an agent's preferences over different strategies can only reveal the margins of her beliefs. However, classical notions of Bayesian updating regard the joint distribution. We develop the relevant environment and tools to solve this issue: We introduce a necessary and sufficient condition on the margins of an agent's beliefs to be consistent with an exchangeable process. Such a consistent process is typically not unique; contemporaneous correlation cannot be identified. We conclude that contemporaneous correlations do not affect the optimal strategy in classical bandit problems.

© 2019 Published by Elsevier Inc.

1. Introduction

This paper studies the structure of beliefs revealed by the preferences of an economic agent who chooses a single action out of several in every period, and determines under what conditions these beliefs are compatible with an exchangeable process.

A leading example for such an economic environment are exploration problems,¹ where an agent has to choose, every period, one project out of several. By observing the outcome of a chosen project, the agent learns both about the chosen project and, in case the outcomes across different projects are correlated, about other projects as well. Therefore, the agent's optimal strategy is a function of the history of observed outcomes, and her beliefs regarding the true process generating the outcomes of each project.

While the generating process jointly determines all projects' outcomes each period, when considering a choice strategy, an expected utility maximizing agent cares only about the outcome of the chosen project. As such, the agent's behavior in exploration problems can reveal only the margins of her beliefs—her beliefs about each individual project conditional on the history.² Determining whether these marginals are “consistent”—if the agent's beliefs at different times are related via Bayesian updating and information arrival—requires a richer understanding of the agent's beliefs. Classical statistical tools

[☆] The authors wish to thank Eddie Dekel, Shaowei Ke, Ehud Lehrer, Kyoungwon Seo, Teddy Seidenfeld, Eran Shmaya, and an anonymous reviewer for their comments and remarks. Special acknowledgment to Kfir Eliaz for discussions which motivated us to think about this project.

* Corresponding author.

E-mail addresses: evan.piermont@rhul.ac.uk (E. Piermont), rteper@pitt.edu (R. Teper).

¹ The modern treatment of exploration models were introduced by Robbins (1952), building on earlier ideas of Thompson (1933), and have since been extensively studied in the statistics literature (as bandit problems), and widely incorporated in economic models (as search problems, stopping problems, research and development, experimentation, portfolio design, etc). See Berry and Fristedt (1985) for an overview of classic results within the statistics literature. For a survey of economic applications see Bergemann and Välimäki (2008).

² In Appendix A, we provide a theoretical framework substantiating this observation.

regard the full joint distribution over the projects' outcomes. Hence, without the analysis that follows, an analyst would not be able to determine if the agent in an exploration problem is statistically sophisticated.

By considering the proper environment and developing the relevant tools we answer the following three questions: (i) What restrictions on the marginal beliefs ensure they are consistent with an exchangeable process jointly determining the projects' outcomes? (ii) When the marginals are consistent with an exchangeable process, is such a process unique? and (iii) Can we draw insights from our identification on the theory of exploration problems?

Recall, an exchangeable process is one in which the belief does not depend on the order of information arrival. Exchangeability has long been the cornerstone of the subjectivist, Bayesian paradigm in the context of repeated experimentation,³ and our interest in exchangeability is tantamount to an assumption that the agent places no special importance to the period in which an outcome was observed. Note, however, that de Finetti's exchangeability condition can not be directly tested in our framework, since the agent does not observe the outcomes of different projects simultaneously. Thus, to answer (i), we provide a condition termed *Across-Marginal Exchangeability* (AM-EXCH), which dictates that the marginal beliefs are invariant to jointly permuting both the order in which projects are chosen and the corresponding outcomes. Across-Marginal Exchangeability is clearly necessary for the agent's beliefs (i.e., the marginal beliefs assessable to the modeler) to coincide with the marginals of an exchangeable process. We here show that it is also sufficient, generalizing de Finetti's representation result to frameworks (such as the exploration environment) in which only marginals can be observed.

The answer to (ii) is more subtle. In the subsequent Section 2, we show that in the finite horizon case, where the agent chooses strategies over n -periods, her beliefs are sometimes (but not always) uniquely identified. When considering the infinite horizon, the identification problem is more severe despite the fact that the modeler has access to more data. This is because Across-Marginal Exchangeability imparts more constraints as the horizon is longer, excluding the cases where identification is possible. For infinite exploration problems, the consistent exchangeable model is *never* completely identified. In particular, *contemporaneous correlation* (i.e., the likelihood of an outcome of project a in a period given the outcome of project b in the *same* period) carries no economic content in such exploration problems. Nevertheless, we can still point to a meaningful representative even under the partial identification: within the class of exchangeable processes consistent with the agent's marginal beliefs there is a unique process for which outcomes are contemporaneously independent.

And so, to answer (iii), the optimal strategy in infinite horizon bandit problems do not depend on contemporaneous correlations if the ex-ante description of the problem is exchangeable. While this is a negative result from the modeler's vantage—the general stochastic process governing beliefs can only be partially identified—it is a boon to the agent: when solving an exploration problem, contemporaneous correlations can be ignored without changing the set of optimal strategies, simplifying her decision problem.

2. Beliefs and the value of strategies

Consider a standard exploration problem. There is a finite set of consumption outcomes X , over which a utility function $u : X \rightarrow \mathbb{R}$ is defined, and a compact, metrizable set of actions \mathcal{A} , where each action $a \in \mathcal{A}$ can yield any of the outcomes in $S_a \subseteq X$. Each period the agent has to choose one (and only one) action, the outcomes of which she observes and derives utility from. A history of length $n \in \mathbb{N}$ is a sequence of n action-outcome pairs, $h = (a_1, x_1; a_2, x_2; \dots; a_n, x_n)$, where $x_i \in S_{a_i}$ for every $i \in n$. The set of finite histories is denoted by \mathcal{H} . Similarly, an infinite history is an infinite sequence of action-outcome pairs, $(a_1, x_1; a_2, x_2; \dots; a_i, x_i; \dots)$, where $x_i \in S_{a_i}$ for every $i \in \mathbb{N}$. Future payoffs are discounted by $\delta \in (0, 1)$, thus an infinite history $\hat{h} = (a_1, x_1; a_2, x_2; \dots)$ is valued according to discounted utility,

$$U(\hat{h}) = \sum_{i \in \mathbb{N}} \delta^{i-1} u(x_i).$$

A strategy in this environment is a function, $\sigma : \mathcal{H} \rightarrow \mathcal{A}$, determining which action to take following every possible finite history. Since the outcome of a given action is uncertain, the agent's beliefs determine which action she prefers to take following every history, and in sum, her optimal strategy. Towards formalizing this, let $\mathcal{S}_{\mathcal{A}} \equiv \prod_{a \in \mathcal{A}} S_a$, and $\mathcal{S} \equiv \prod_{n \geq 1} \mathcal{S}_{\mathcal{A}}$ (when considering a finite horizon problem, \mathcal{S} refers to the n -fold product of $\mathcal{S}_{\mathcal{A}}$). The set \mathcal{S} represents the grand state-space; a state determines the realization of each action in each period. The uncertainty over the state space, that is the agent's beliefs over what is the state generating the actions' outcomes, is typically captured in applications through a probability $\zeta \in \Delta(\mathcal{S})$.

Given such a belief, ζ , every strategy, σ , induces a unique countably additive probability \mathbf{P}_{σ} over the set of infinite histories.⁴ The agent values a strategy according to its expected utility with respect to the probability it induces over infinite histories. That is,

$$V(\sigma) = \mathbf{E}_{\sigma} \left(U(\hat{h}) \right), \quad (1)$$

³ See for example de Finetti (1972); Persi (1977); Schervish (2012).

⁴ Endowed with the Borel sigma-algebra generated by all finite histories. We identify each finite history with the set of infinite histories that extend it.




Fig. 1. Figure 1A:J shows the joint distribution for Example 1.A and 1B:J shows the joint distribution for Example 1.B.

where E_σ denotes the expectation operator with respect to \mathbf{P}_σ . The agent's optimal strategy, if such exists, is the one maximizing $V(\cdot)$.

Remark 1. Denote by $\mu_{h,a}(x)$ the ζ -probability, that conditional on an n -period history h , action a yields outcome x in period $n + 1$. Given two agents (U, ζ) and (U, ζ') such that $\{\mu_{h,a}\}_{h \in \mathcal{H}, a \in \mathcal{A}} = \{\mu'_{h,a}\}_{h \in \mathcal{H}, a \in \mathcal{A}}$ then $V = V'$: the agents rank strategies identically.

In other words, while the probability ζ completely specifies the underlying uncertainty over the joint realization of all actions following every history, the agent's valuations are determined entirely by the marginals of this process: $\{\mu_{h,a}\}_{h \in \mathcal{H}, a \in \mathcal{A}}$. To obtain Remark 1, notice that for a finite history, $h \in \mathcal{H}$, and an action, $a \in \mathcal{A}$, $\mathbf{P}_\sigma(h; a, x)$ is defined by

$$\mathbf{P}_\sigma(h; a, x) = \mathbf{P}_\sigma(h)\mu_{h,a}(x) \tag{2}$$

if $\sigma(h) = a$ and $x \in S_a$. Otherwise, the probability is 0. By standard arguments \mathbf{P}_σ is uniquely determined by its measure of finite histories. By examining (2) it is clear that \mathbf{P}_σ , and therefore E_σ , depends only on $\{\mu_{h,a}\}_{h \in \mathcal{H}, a \in \mathcal{A}}$.

Remark 1 indicates that an analyst with access to data regarding choice or preference over strategies in an exploration problem—no matter how detailed—can never identify more than the marginals of the agent's beliefs. The remainder of this paper explores the limits of inference that can be made when the marginals, but nothing more, are identified. Of course, Remark 1 does not ensure that $\{\mu_{h,a}\}_{h \in \mathcal{H}, a \in \mathcal{A}}$ are themselves identifiable from any type of observable data. This latter question is formally answered in the affirmative in the supplemental appendix—we conduct a decision theoretic exercise, construct the set of exploration strategies, and provide the axioms allowing us to determine whether the agent is indeed a discounted subjective expected utility maximizer, as in Eq. (1). We uniquely identify the marginals, $\{\mu_{h,a}\}_{h \in \mathcal{H}, a \in \mathcal{A}}$ (and the utility parameters, u and δ).

Example 1.A. Consider a two-period problem where in each period the agent has to choose between two projects, a and b , each of which can either succeed or fail: $S_a = \{s_a, f_a\}$ and $S_b = \{s_b, f_b\}$. The agent believes that each project will have exactly one success, equally likely to be in either period, and, moreover, believes the two projects will succeed and fail jointly.⁵

The corresponding process $\zeta \in \Delta(\mathcal{S})$ is given by the table in Fig. 1(1A:J); a row corresponds to a joint outcome of the two projects in period $n = 1$, and a column to a joint outcome of the two projects in period $n = 2$. By Remark 1, we assume the modeler cannot observe ζ itself but instead the marginals,

$$\mu_{\emptyset}(h_x) = \frac{1}{2} \quad \mu_{(x,h_x)}(g_y) = 1 \quad \mu_{(x,h_x)}(h_y) = 0. \tag{1A:M}$$

Where $x, y \in \{a, b\}$ and $h, g \in \{s, f\}$ with $h \neq g$.

Assume further that the per-period utility associated with each outcome is $u(s_a) = 9$, $u(f_a) = -9$, $u(s_b) = 18$, and $u(f_b) = -18$. The agent is an expected utility maximizer, and her total utility is the sum across the two periods. Given these restrictions on preferences, (1A:M) determines the agent's valuation of strategies. Indeed, for $x, y, z \in \{a, b\}$, let $(x, (y, z))$ denote the strategy in which the agent takes action x in the first period, and y conditional on x 's success and z on x 's failure. The agent's valuations for strategies are given as follows: $V(x, (y, z)) = 0$ if $y = z$, and

$$\begin{aligned} V(a, (a, b)) &= V(b, (a, b)) = \frac{9}{2} \\ V(a, (b, a)) &= V(b, (b, a)) = -\frac{9}{2}. \end{aligned} \tag{1A:P}$$

⁵ Examples 1.A and 1.B would have similar implications if we considered a somewhat less extreme point of view in terms of the probabilities. For instance, the same conclusions would have been reached had we considered projects yielding s or f in the first period with equal probability, while in the second period, a project that was a success in the first period also yields s in the second period with probability $\frac{4}{9}$, and given a failure in the first period, a project yields s in the second with probability $\frac{5}{9}$. Such negative auto-correlated processes are the typical representative of exchangeability in finite horizon models, and have been used extensively in the finance literature (see for example Poterba and Summers (1988); Berk and Green (2004) and references therein).

		n = 2						n = 2			
		s _a , s _b	s _a , f _b	f _a , s _b	f _a , f _b			s _a , s _b	s _a , f _b	f _a , s _b	f _a , f _b
n = 1	s _a , s _b	5/16	0	0	3/16	n = 1	s _a , s _b	41/256	15/256	15/256	9/256
	s _a , f _b	0	0	0	0		s _a , f _b	15/256	9/256	9/256	15/256
	f _a , s _b	0	0	0	0		f _a , s _b	15/256	9/256	9/256	15/256
	f _a , f _b	3/16	0	0	5/16		f _a , f _b	9/256	15/256	15/256	41/256

Fig. 2. Two alternative joint distributions discussed in Example 2.

The exchangeable belief, ζ as defined by Fig. 1(1A:J), is in fact uniquely determined by the marginal beliefs. Although we do not assume that ζ is observed, it is identified by the agent’s preferences over strategies. In particular, we can identify that the agent believes the two projects are perfectly correlated.

Example 1.B. If, instead, the agent believed each project will have exactly one success, equally likely to be in either period (as above), but unlike the previous example believed that the projects were independent of each other, she would entertain the joint distribution given by Fig. 1(1B:J), the marginals of which are,

$$\begin{aligned}
 \mu_{\emptyset}(h_x) &= \frac{1}{2} & \mu_{(x,h_x)}(g_x) &= 1 & \mu_{(x,h_x)}(h_x) &= 0 \\
 \mu_{(x,h_x)}(g_y) &= \frac{1}{2} & \mu_{(x,h_x)}(h_y) &= \frac{1}{2}.
 \end{aligned}
 \tag{1B:M}$$

Where $x, y \in \{a, b\}$ with $x \neq y$ and $h, g \in \{s, f\}$ with $h \neq g$. This of course has a corresponding change in the agent’s valuations: $V(x, (y, z)) = 0$ if $y = z$, and

$$\begin{aligned}
 V(a, (a, b)) &= -\frac{9}{2}V(b, (a, b)) = 9 \\
 V(a, (b, a)) &= \frac{9}{2}V(b, (b, a)) = -9.
 \end{aligned}
 \tag{1B:P}$$

While the above example shows that the correlation between projects can potentially affect (or, be recovered from) the agent’s preferences—or equivalently her marginal beliefs—it is not typical. The inherent observability constraint in the standard framework of experimentation generally bears a cost; the exchangeable process with which our observables are consistent is often non-unique.

Example 2. Let the actions and outcomes be the same as Example 1.A. The agent considers two equally probable possibilities: in the first both projects have a $\frac{1}{4}$ likelihood of succeeding in both periods (i.e., i.i.d. over time, with probability $\frac{1}{4}$) and in the second the likelihood of success is $\frac{3}{4}$. Consider the two joint distributions in Fig. 2. The left panel is the joint distribution when the agent believes the two projects intra-period successes and failures are perfectly correlated, whereas the right is when they are perfectly independent.

Under the case of perfect correlation, the outcome of project a in period 0 perfectly identifies what would have happened had project b been chosen instead. At first glance, this information seems valuable for the agent’s exploration problem, however, both joint distributions impart the exact same restrictions on marginal beliefs:

$$\mu_{\emptyset}(h_x) = \frac{1}{2} \quad \mu_{(x,h_x)}(g_x) = \frac{3}{16} \quad \mu_{(x,h_x)}(h_x) = \frac{5}{16}
 \tag{2:M}$$

Where $x, y \in \{a, b\}$ and $h, g \in \{s, f\}$ with $h \neq g$. Therefore, the agent’s valuation of all strategies, and in particular her optimal strategy, is unaffected by the correlation between the two actions.

In Example 1 the agent’s preferences over strategies perfectly revealed her perceived *contemporaneous* correlation between the two projects. In Example 2, we can infer nothing about how the agent perceives the contemporaneous correlation. The latter proves to be the rule. In the sequel, we show that in the infinite horizon problem, beliefs can *never* be fully identified. Fortunately, the obstruction can be precisely delineated; contemporaneous correlations stand as the only obstacle thwarting the identification of the agent’s joint beliefs.

3. The statistical framework

In order for a modeler to understand the DM’s updating process, and whether it follows Bayes rule, we need to construct her beliefs regarding not only each action individually but also her beliefs regarding the possibly joint outcomes across all actions (in particular, the correlation between actions). As we will see, in the generic case we have insufficient data to uniquely identify a (subjective) joint distribution. We will still, however, be able to identify a representative with unique properties.

Observable Processes. Consider the family \mathcal{T} of all sequences of individual experiments (i.e., individual actions), where different experiments can be taken in the different periods. Let $\mathbf{T} = (T_1, T_2, \dots)$ where $T_i \in \{S_a : a \in \mathcal{A}\}$ for every $i \geq 1$; so, each T_i corresponds to taking an action, say a , and expecting one of its outcomes, S_a . (Like before S_a corresponds to the set of possible outcomes.) \mathcal{T} is then the collection of all such \mathbf{T} 's. For each $\mathbf{T} = (S_{a_1}, S_{a_2}, \dots)$ let $\zeta_{\mathbf{T}} \in \Delta^{\mathcal{B}}(\mathbf{T})$ be a process over \mathbf{T} ; a distribution over all possible outcomes when taking action a_1 , followed by a_2 , followed by a_3 , etc.⁶ For a given history of action-outcomes pairs, $h \in (a_1, x_1 \dots a_n, x_n)$, we say $h \in \mathbf{T} = (T_1, \dots, T_n, T_{n+1}, \dots)$ whenever $S_{a_i} = T_i$ for $1 \leq i \leq n$ (and maintaining the above assumption that $x_i \in S_{a_i}$). In a slight abuse of notation, we can identify each $h = (a_1, x_1 \dots a_n, x_n) \in \mathbf{T}$ with the event in \mathbf{T} such that first n outcomes are $(x_1 \dots x_n)$, so as to make sense of $\zeta_{\mathbf{T}}(h)$ and $\zeta_{\mathbf{T}}(\cdot | h)$. Lastly, for a sequence of experiments $\mathbf{T} = (T_1, \dots, T_n, T_{n+1}, \dots)$ and a permutation $\pi : \{1, \dots, n\} \rightarrow \{1, \dots, n\}$, denote $\pi\mathbf{T} = (T_{\pi(1)}, \dots, T_{\pi(n)}, T_{n+1}, \dots)$.

A **Subjective Expected Experimentation (SEE) belief structure** is a family of processes $\{\zeta_{\mathbf{T}}\}_{\mathbf{T} \in \mathcal{T}}$, such that for any $\mathbf{T}, \mathbf{T}' \in \mathcal{T}$ and $h \in \mathcal{H}$ if $h \in \mathbf{T}$ and $h \in \mathbf{T}'$, then $\zeta_{\mathbf{T}}(h) = \zeta_{\mathbf{T}'}(h)$. This condition imposes that the probability of outcomes today do not depend on which experiments are to be conducted in the future.

Remark 1 in the previous section shows that the DM's belief over S , ζ , is identified only up to the marginal beliefs, $\{\mu_{h,a}\}_{h \in \mathcal{H}, a \in \mathcal{A}}$ (and this identification is tight, as shown in the appendix). Each such family of marginals uniquely determines an SEE belief structure in the obvious manner. Given $\{\mu_{h,a}\}_{h \in \mathcal{H}, a \in \mathcal{A}}$ and a sequence $\mathbf{T} = (S_{a_1}, S_{a_2}, \dots)$, $\zeta_{\mathbf{T}}$ is the unique (countably additive) process satisfying

$$\zeta_{\mathbf{T}}(h) = \mu_{\emptyset, a_1}(x_1) \cdot \mu_{(a_1, x_1), a_2}(x_2) \cdots \mu_{(a_1, x_1, \dots, a_{n-1}, x_{n-1}), a_n}(x_n)$$

for all $h \in \mathcal{H}$. In fact, SEE belief structures are exactly the set of processes that can be constructed from a family of marginal beliefs, $\{\mu_{h,a}\}_{h \in \mathcal{H}, a \in \mathcal{A}}$.

Exchangeable Processes and Consistency. Recall, $S_{\mathcal{A}} \equiv \prod_{a \in \mathcal{A}} S_a$, and $S \equiv \prod_{n \geq 0} S_{\mathcal{A}}$. S represents the grand state-space; a state determines the realization of each action in each period.

Now, we say that an SEE belief structure $\{\zeta_{\mathbf{T}}\}_{\mathbf{T} \in \mathcal{T}}$ is **consistent** with $\zeta \in \Delta^{\mathcal{B}}(S)$ if $\zeta|_{\mathbf{T}}$ (that is, the marginal of ζ to \mathbf{T}) coincides with $\zeta_{\mathbf{T}}$, for every $\mathbf{T} \in \mathcal{T}$. In such a case the processes ζ , which we cannot observe, explains all our data.

Because it forms the basis subjective Bayesianism and for the statistical literature on exploration problems, we will pay particular attention to the class of *exchangeable* processes.

Definition. Let Ω be a probability space and $\hat{\Omega} = \prod_{n \geq 1} \Omega$. The process $\zeta \in \Delta^{\mathcal{B}}(\hat{\Omega})$ is **exchangeable** if for any finite permutation $\pi : \mathbb{N} \rightarrow \mathbb{N}$ and event $E = \prod_{n \in \mathbb{N}} E_n$, we have

$$\zeta(E) = \zeta\left(\prod_{n \in \mathbb{N}} E_{\pi(n)}\right). \tag{3}$$

From the economic vantage, a DM who understands there to be an exchangeable process governing the outcome of actions would be considered Bayesian. A DM with exchangeable beliefs (acts as if she) entertains a prior on the data generating parameter and updates her beliefs following every observation. This interpretation is due to the fundamental results of de Finetti (1931, 1937) and later extensions of Hewitt and Savage (1955).

Remark 2 (de Finetti). The process $\zeta \in \Delta^{\mathcal{B}}(\hat{\Omega})$ is exchangeable if and only if there exists a probability measure θ over $\Delta^{\mathcal{B}}(\Omega)$, such that

$$\zeta(E) = \int_{\Delta^{\mathcal{B}}(\Omega)} \hat{D}(E) d\theta(D), \tag{4}$$

where for any $D \in \Delta^{\mathcal{B}}(\Omega)$, \hat{D} is the corresponding product measure over $\hat{\Omega}$. Moreover, θ is unique.

We would like to understand under what circumstances an SEE belief structure is a result of Bayesian updating. If we could infer from preferences the beliefs over the joint realizations of all actions, that is $\prod_{a \in \mathcal{A}} S_a$, then our questions would boil down to verifying whether this process satisfies exchangeability. However, we can only infer the beliefs over each action separately, and thus, our task remains. We need to find a condition on the family of $\zeta_{\mathbf{T}}$'s that determines whether it follows Bayes rule.

Definition. An SEE belief structure $\{\zeta_{\mathbf{T}}\}_{\mathbf{T} \in \mathcal{T}}$ is **Across-Marginal Exchangeability (AM-EXCH)** if

$$\zeta_{\mathbf{T}}(h) = \zeta_{\pi\mathbf{T}}(\pi h)$$

for every $\mathbf{T} \in \mathcal{T}$, $h \in \mathbf{T}$ and a finite permutation $\pi : \mathbb{N} \rightarrow \mathbb{N}$.

⁶ For any metric space M , denote $\Delta^{\mathcal{B}}(M)$ as the set of Borel probability distributions over M , endowed with the weak*-topology.

Intuitively, AM-EXCH requires that if we consider a different order of experiments, then the probability of outcomes (in the appropriate order) does not change. The next theorem states that Across-Marginal Exchangeability is a necessary and sufficient condition for an SEE belief structure to be consistent with Bayesian updating of some belief over the joint realizations of all actions.

Theorem 1. An SEE belief structure $\{\zeta_{\mathbf{T}}\}_{\mathbf{T} \in \mathcal{T}}$ satisfies AM-EXCH if and only if it is consistent with an exchangeable process $\zeta \in \Delta^{\mathcal{B}}(\mathcal{S})$.

Theorem 1 is stated without proof. Necessity is trivial and sufficiency will be a straightforward application of Theorem 2. Although the theorem as stated concerns only infinite-horizon processes, AM-EXCH is also a necessary and sufficient condition for a finite horizon process to be consistent with some exchangeable process, provided there exists some consistent joint distribution.⁷

4. Strong exchangeability

In this section we introduce a strengthening of exchangeability, aptly called *strong exchangeability*, that corresponds to the maximal preservation of symmetry implied by AM-EXCH. Strongly exchangeable processes are those under which each dimension can be permuted independently. It will turn out, in the infinite horizon, strongly exchangeable processes can be characterized as those in which stochastic independence is preserved both inter-temporally (as in vanilla exchangeability) and contemporaneously.⁸ Putting these results together: if a modeler can only observe the marginals of a DM's beliefs, and those marginals are consistent with any exchangeable process (i.e., satisfy AM-EXCH), then the modeler identifies nothing about the DM's perceived correlation between projects.

Definition. A process $\zeta \in \Delta^{\mathcal{B}}(\mathcal{S})$ is **strongly exchangeable** if for any set of finite permutations $\{\pi_a : \mathbb{N} \rightarrow \mathbb{N}\}_{a \in \mathcal{A}}$ and event $E = \prod_{n \in \mathbb{N}} \prod_{a \in \mathcal{A}} E_{n,a}$, we have

$$\zeta(E) = \zeta\left(\prod_{n \in \mathbb{N}} \prod_{a \in \mathcal{A}} E_{\pi_a(n),a}\right). \tag{5}$$

Setting $\pi_a = \pi_b$ for all $a, b \in \mathcal{A}$, delivers the definition of exchangeability. Following the intuition above, it should come as no surprise that under AM-EXCH, strong exchangeability can never be ruled out. In other words, there is no SEE belief structure—therefore no preference over exploration problems—that distinguishes exchangeability from strong exchangeability. Strongly exchangeable processes are especially relevant with respect to the current focus because they act as representative members to the equivalence classes of exchangeable processes consistent with the same SEE belief structure.

Theorem 2. An SEE belief structure $\{\zeta_{\mathbf{T}}\}_{\mathbf{T} \in \mathcal{T}}$ satisfies AM-EXCH if and only if it is consistent with a strongly exchangeable process. Furthermore, such a strongly exchangeable process is unique.

Proof. In Section 6. ■

While there are obvious conceptual similarities, Theorem 2 (and by proxy, Theorem 1) do not straightforwardly follow from the extant results regarding exchangeability. Because only the marginals of the DM's beliefs are observable, the standard definitions of symmetry (or other characterizations of exchangeability) cannot be directly applied. The proof of Theorem 2 explicitly constructs a consistent, strongly exchangeable process.

Briefly: consider an event, $E \subset \mathcal{S}$ for which $E_n = \mathcal{S}_a$ for all sufficiently large n — E only places restrictions on the observations for a finite number of periods. The set of all such events generate the relevant σ -algebra over \mathcal{S} , and so determining the value of a process on all such events uniquely determines the process. For any such event, E , we can permute E in each dimension to construct another event \hat{E} , for which at most one action is restricted at any given time (that is $\{a \in \mathcal{A} \mid \hat{E}_{n,a} \neq \mathcal{S}_a\}$ has at most one element for each n). Since only one action is restricted per period, the probability of \hat{E} can be identified by the SEE belief structure, and, under the assumption of strong exchangeability, we know this also determines the probability of E . The proof shows this process is well defined.

Just as infinite horizon exchangeable processes can be characterized as being a mixture of i.i.d. distributions, infinite horizon strongly exchangeable are exactly mixtures of distributions that are both inter-temporally i.i.d. and lack any correlation between different projects within the same period.

⁷ The proof in the finite horizon case is quite intuitive. Let η be a consistent joint distribution. For each event E let E^* denote the union of πE for all $n!$ permutations $\pi : n \rightarrow n$, where n is the number of periods. Construct ζ as follows: $\zeta(E) = \frac{\eta(E^*)}{n!}$. The process ζ is well defined and it is clearly exchangeable. Moreover, $\{\zeta_{\mathbf{T}}\}_{\mathbf{T} \in \mathcal{T}} = \{\eta_{\mathbf{T}}\}_{\mathbf{T} \in \mathcal{T}}$, since for all $E \in \mathcal{T}$, $\eta(E)$ is equal to $\eta(\pi E)$ and therefore also to $\zeta(E)$.

⁸ We feel reasonably certain that strong exchangeability must have been studied previously in the statistics literature. However, we have found no references.

Theorem 3. The process $\zeta \in \Delta^{\mathcal{B}}(\mathcal{S})$ is strongly exchangeable if and only if there exists a probability measure θ over $\Delta^{IN} \equiv \prod_{a \in \mathcal{A}} \Delta(S_a)$, such that

$$\zeta(E) = \int_{\Delta^{IN}} \hat{D}(E) d\theta(D),$$

where for any $D \in \Delta^{IN}$, \hat{D} is the corresponding product measure over \mathcal{S} .

Proof. In Section 6. ■

Under a strongly exchangeable process, the DM believes the outcomes of actions that occur at the same time are independently resolved. Of course, this does not impose that there is no informational cross contamination between actions. Information regarding the distribution of action a is informative about the underlying parameter governing the exchangeable process, and therefore, also about the distribution of action b .

5. A comment on Bayesianism in environments of experimentation

The results in Section 4 have two related implications to Bayesianism in general models of experimentation. First, it is well known that the beliefs of two Bayesians observing the same sequence of signals will converge in the limit. Our results imply that in a setup of experimentation, different Bayesians obtaining the same information, might still hold different views of the world in the limit. Their beliefs over the uncertainty underlying each action will be identical, but they can hold different beliefs over the joint distribution.

The second point has to do with the possible equivalence with non-Bayesian DMs. Theorem 2 states that AM-EXCH is necessary and sufficient for an SEE belief system to be consistent with some exchangeable process—but it might also be consistent with a non-exchangeable process. Consider the following example of a stochastic process. In every period two coins are flipped. In odd periods the coins are perfectly correlated (with equal probability on HH and TT), and in even periods the coins are identical and independent (and both have equal probability on H and T). The associated observable processes satisfy AM-EXCH, but the process itself is clearly not exchangeable. Nevertheless, Theorem 2 guarantees that there is a (unique) strongly exchangeable process that is consistent with the SEE belief structure.

6. Proofs

Proof of Theorem 2. Fix an SEE belief structure $\{\zeta_{\mathbf{T}}\}_{\mathbf{T} \in \mathcal{T}}$. We first construct a pre-measure $\hat{\zeta}$ on the semi-algebra of cylinder sets. Fix any well-ordering over \mathcal{A} . Set $\hat{\zeta}(\emptyset) = 0$ and $\hat{\zeta}(\mathcal{S}) = 1$. Let $E \neq \mathcal{S}$ be an arbitrary cylinder, i.e., $E = \prod_{n \in \mathbb{N}} \prod_{a \in \mathcal{A}} E_{n,a}$, such that for only finitely many (n, a) , is $E_{n,a} \neq S_a$. Clearly, there are a finite number of $a \in \mathcal{A}$ such that $E_{k,a} \neq S_a$ for any k . By the ordering on \mathcal{A} denote these $a_1 \dots a_n$. For each a_i let m_i denote the number of components such that $E_{k,a_i} \neq S_{a_i}$, and for $j = 1 \dots m_i$, let $k_{i,j}$ denote the j^{th} such component. Finally, for each a_i , let π_{a_i} denote any permutation such that $\pi_{a_i}(k_{i,j}) = j + \sum_{i' < i} m_{i'}$. Consider $\hat{E} = \prod_{n \in \mathbb{N}} \prod_{a \in \mathcal{A}} E_{\pi_a(n),a}$, where $\pi_a = \pi_{a_i}$ if $a \in a_1 \dots a_n$ and the identity otherwise. That is, for $n \in 1 \dots m_1$, $\hat{E}_{n,a} = S_a$ for all a except a_1 , for $n \in m_1 + 1 \dots m_1 + m_2$, $\hat{E}_{n,a} = S_a$ for all a except a_2 , etc. Let $\mathbf{T}(E)$ denote any sequence such that $T_n = S_{a_i}$ for $\sum_{i' < i} m_{i'} < n \leq \sum_{i' \leq i} m_{i'}$. Again that is, for $n \in 1 \dots m_1$, $T_n = S_{a_1}$, for $n \in m_1 + 1 \dots m_1 + m_2$, $T_n = S_{a_2}$, etc.

For the remainder of this proof, for any cylinder E , \hat{E} denotes the corresponding cylinder generated by the above process, in which at most a single action is restricted in each period. Let $\mathbf{T}(E)$ denote any observable process which observes the sequence of restricted actions. Finally, for any cylinder, E , which is restricted in most one action each period, and any \mathbf{T} which observes each restricted set, identify E the relevant event in \mathbf{T} . So, set $\hat{\zeta}(E) = \zeta_{\mathbf{T}(E)}(\hat{E})$. This is well defined by the restriction of SEE belief structures.

To apply the Carathéodory extension theorem for semi-algebras, we need to show that for any sequence of disjoint cylinders $\{E^k\}_{k \in \mathbb{N}}$ such that $E = \bigcup_{k \in \mathbb{N}} E^k$ is a cylinder, $\hat{\zeta}(E) = \sum_{k \in \mathbb{N}} \hat{\zeta}(E^k)$. Towards this, assume that E, E' are disjoint cylinders such that $E \cup E'$ is a cylinder. Then it must be that there exists a unique (n, a) such that $E_{n,a} \cap E'_{n,a} = \emptyset$ and for all other (m, b) , $E_{m,b} = E'_{m,b}$. Indeed, if this was not the case, then there exists some (m, b) and some x such that (WLOG) $x \in E_{m,b} \setminus E'_{m,b}$. But then, for all $s \in E \cup E'$, $s_{m,b} = x \implies s_{n,a} \in E_{n,a} \neq (E \cup E')_{n,a}$ a contradiction to $E \cup E'$ being a cylinder. But this implies \hat{E} and \hat{E}' induce the same sequence of restricted coordinates, differing on the restriction of single coordinate, and therefore, $\mathbf{T}(E) = \mathbf{T}(E')$. This implies that $\hat{E} \cup \hat{E}' \subseteq \mathbf{T}(E)$. Since $\zeta_{\mathbf{T}(E)}$ is finitely additive, so therefore $\hat{\zeta}(E \cup E') = \zeta_{\mathbf{T}(E)}(\hat{E} \cup \hat{E}') = \zeta_{\mathbf{T}(E)}(\hat{E}) + \zeta_{\mathbf{T}(E)}(\hat{E}') = \hat{\zeta}(E) + \hat{\zeta}(E')$.

Since $\hat{\zeta}$ is finitely additive over cylinder sets, countable additivity follows if we show that for all decreasing sequences of cylinders $\{E^k\}_{k \in \mathbb{N}}$, such that $\inf_k \hat{\zeta}(E^k) = \epsilon > 0$, we have $\bigcap_{k \in \mathbb{N}} E^k \neq \emptyset$. But this follows immediately from the finiteness of S_a . Since $E^{k+1} \subseteq E^k$, it must be that $E^k_{n,a} \subseteq E^{k+1}_{n,a}$. But each $E^k_{n,a}$ is finite, hence compact, and nonempty, because $\zeta(E^k) \geq \epsilon$. Therefore $\bigcap_{k \in \mathbb{N}} E^k_{n,a} \neq \emptyset$. The result follows by noting that the intersection of cylinder sets is the cylinder generated by the intersection of the respective generating sets. Let ζ denote the unique extension of $\hat{\zeta}$ to the σ -algebra on \mathcal{S} .

That ζ is consistent with $\{\zeta_{\mathbf{T}}\}_{\mathbf{T} \in \mathcal{T}}$ is immediate. We need to show that ζ is strongly exchangeable. Let E be a cylinder. Let $\tilde{\pi}_a$ denote a finite permutation for each $a \in \mathcal{A}$. Let $F = \prod_{n \in \mathbb{N}} \prod_{a \in \mathcal{A}} E_{\tilde{\pi}_a(n), a}$. Let π_{a_i} denote the permutation given by the construction of \hat{F} . Then $\hat{F} = \prod_{n \in \mathbb{N}} \prod_{a \in \mathcal{A}} E_{\pi_a(\tilde{\pi}_a(n)), a}$. This implies there exists some permutation π^* such that $\hat{F} = \prod_{n \in \mathbb{N}} \prod_{a \in \mathcal{A}} E_{\pi^*(n), a}$. By AM-EXCH, $\zeta_{\mathbf{T}(\hat{E})}(\hat{E}) = \zeta_{\pi^* \mathbf{T}(\hat{E})}(\pi^* \hat{E}) = \zeta_{\mathbf{T}(\hat{F})}(\hat{F})$. Therefore, $\zeta(E) = \zeta(F)$ and so, ζ is strongly exchangeable.

Finally, the similar logic show that ζ is unique. Towards a contradiction, assume there was some distinct, strongly exchangeable ζ' , also consistent with $\{\zeta_{\mathbf{T}}\}_{\mathbf{T} \in \mathcal{T}}$. Then, since the cylinder sets form a π -system, there must be some cylinder such that $\zeta(E) \neq \zeta'(E)$. But, by strong exchangeability, $\zeta(\hat{E}) = \zeta(E)$ and $\zeta'(\hat{E}) = \zeta'(E)$, so $\zeta(\hat{E}) \neq \zeta'(\hat{E})$ – a contradiction to their joint consistency with $\{\zeta_{\mathbf{T}}\}_{\mathbf{T} \in \mathcal{T}}$. ■

Proof of Theorem 3. First we show that if a process ζ over \mathcal{S} is both i.i.d. (with marginal $D \in \Delta(\mathcal{S}_{\mathcal{A}})$) and strongly exchangeable, then it must be that the marginals of D (on $\{S_a\}_{a \in \mathcal{A}}$) are independent, that is $D \in \Delta^{IN}$. Indeed, consider two non-empty, disjoint collection of actions, $\hat{\mathcal{A}}, \hat{\mathcal{A}}' \subset \mathcal{A}$. Let $E, F \in \mathcal{S}_{\hat{\mathcal{A}}}, E', F' \in \mathcal{S}_{\hat{\mathcal{A}}'}$, be measurable events. Identify E^n with the cylinder in \mathcal{S} : $E^n = \{s \in \mathcal{S} | s_n \in E\}$. Since ζ is strongly exchangeable we have that

$$\zeta(E^n \cap E'^n \cap F^{n+1} \cap F'^{n+1}) = \zeta(E^n \cap F'^n \cap F^{n+1} \cap E'^{n+1}). \tag{2SYM}$$

We will refer to the latter weaker property as *two symmetry*. Now, since ζ is i.i.d. generated by D , we have that (abusing notation by identifying E with the cylinder it generates in $\mathcal{S}_{\mathcal{A}}$)

$$D(E \cap E') \cdot D(F \cap F') = D(E \cap F') \cdot D(F \cap E').$$

Substituting via the rule of conditional probability:

$$D(E|E') \cdot D(E') \cdot D(F|F') \cdot D(F') = D(E|F') \cdot D(F') \cdot D(F|E') \cdot D(E').$$

This implies that

$$\frac{D(E|E')}{D(E|F')} = \frac{D(F|E')}{D(F|F')}.$$

Since this is true for all events, we have that $D(E|E') = D(E|F')$ for every $E \subseteq \mathcal{S}_{\hat{\mathcal{A}}}$ and $E', F' \subseteq \mathcal{S}_{\hat{\mathcal{A}}'}$, implying $\hat{\mathcal{A}}$ and $\hat{\mathcal{A}}'$ are independent.

We now move to show that strong exchangeability is sufficient for the representation specified in the statement of the result. Since strong exchangeability implies exchangeability, we can apply a version de Finetti’s theorem (Schervish (2012) Theorem 1.49) and represent the process ζ by

$$\zeta(\cdot) = \int_{\Delta(\mathcal{S}_{\mathcal{A}})} \hat{D}(\cdot) d\psi(D).$$

We need to show that ψ ’s support lies in Δ^{IN} .

For $s \in \mathcal{S}$ and $t \in \mathbb{N}$ let s_t be the projection of s into the first t periods. Now, let $\zeta(\cdot | s_t) : \mathcal{S}_{\mathcal{A}} \rightarrow [0, 1]$ be the *one period ahead predictive probability*, given that the history of realizations in the first t periods is s_t . Since ζ is exchangeable, $\zeta(\cdot | s_t)$ converges (as $t \rightarrow \infty$) with ζ probability 1. Moreover, the set of all limits is the support of ψ . Denote the limit for a particular s by D_s . Of course, the exchangeability of ζ also guarantees that $\zeta(\cdot, \cdot | s_t) : \mathcal{S}_{\mathcal{A}} \times \mathcal{S}_{\mathcal{A}} \rightarrow [0, 1]$, that is the *two period ahead predictive probability*, converges to $D_s \times D_s$. Furthermore, ζ is strongly exchangeable; the limit itself satisfies (2SYM), and the arguments above imply that $D_s \in \Delta^{IN}$ with ζ probability 1. ■

Appendix. Supplementary material

Supplementary material related to this article can be found online at <https://doi.org/10.1016/j.geb.2019.04.005>.

References

Bergemann, Dirk, Välimäki, Juuso, 2008. Bandit problems. In: Durlauf, Steven N., Blume, Lawrence E. (Eds.), *The New Palgrave Dictionary of Economics*, second edition.

Berk, Jonathan B., Green, Richard C., 2004. Mutual fund flows and performance in rational markets. *J. Polit. Econ.* 112 (6), 1269–1295.

Berry, Donald A., Fristedt, Bert, 1985. *Bandit Problems: Sequential Allocation of Experiments*. Monographs on Statistics and Applied Probability. Springer.

de Finetti, Bruno, 1931. Sul significato soggettivo della probabilità. *Fundam. Math.* 17 (1), 298–329.

de Finetti, Bruno, 1937. La prévision: ses lois logiques, ses sources subjectives. *Ann. Inst. Henri Poincaré* 7, 1–68.

de Finetti, Bruno, 1972. *Probability, Induction, and Statistics*. John Wiley, New York.

Diaconis, Persi, 1977. Finite forms of de finetti’s theorem on exchangeability. *Synthese* 36 (2), 271–281.

Hewitt, Edwin, Savage, Leonard J., 1955. Symmetric measures on Cartesian products. *Trans. Am. Math. Soc.* 80 (2), 470.

Poterba, James M., Summers, Lawrence H., 1988. Mean reversion in stock prices: evidence and implications. *J. Financ. Econ.* 22 (1), 27–59.

Robbins, Herbert, 1952. Some aspects of the sequential design of experiments. *Bull. Am. Math. Soc.* 58 (5), 527–536.

Schervish, Mark J., 2012. *Theory of Statistics*. Springer Science & Business Media.

Thompson, William R., 1933. On the likelihood that one unknown probability exceeds another in view of the evidence of two samples. *Biometrika* 25 (3/4), 285–294.