

Plans of Action*

Evan Piermont[†] and Roe Teper[‡]

October 5, 2016

Abstract

We study the extent to which contemporaneous correlations across actions affect an agent's preferences over the different strategies in exploration problems. We show that such correlations carry no economic content and do not affect the agent's preferences and, in particular, her optimal strategy. We argue that for similar reasons there is an inherent partial identification of the beliefs in exploration problems. Nevertheless, even under the partial identification, we show there are meaningful behavioral restrictions allowing the modeler to test whether the agent is acting according to some Bayesian model.

Key words: Bandit problems; correlated arms; strong exchangeability, responsive learning.

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

[†]University of Pittsburgh, Department of Economics, ehp5@pitt.edu

[‡]University of Pittsburgh, Department of Economics, rteper@pitt.edu

1 INTRODUCTION

Exploration models capture a common trade off between an immediate payoff and new information, which can potentially impact future decisions and payoffs.¹ In such models an agent has to choose, every period, one project out of several in which to invest. By observing the outcome of an investment, the agent learns both about the chosen project and, in case the outcomes across different projects are correlated, about other projects as well. Each decision is predicated on the tradeoff between the immediate value of the investment and the future value of the information obtained by observing the outcome. Therefore, the agent's optimal investment strategy is a function of the history of observed outcomes, the projects that will be feasible in the future, and her beliefs over the true *joint* distribution of the outcomes of each project. While it is the correlation between projects that allows the agent to extrapolate her observations to future outcomes of the different projects, we show in this paper that *contemporaneous correlations* (i.e., the likelihood of an outcome of project a in a period given the outcome of project b in the *same* period) carry no economic content in such exploration problems. In other words, when solving an exploration problem, contemporaneous correlations can be ignored without changing the set of optimal strategies.

Consider, for example, a scenario in which our agent has to invest each period in one of two projects. The outcome of each project depends on the state of the economy, which could be either high or low with equal probabilities, independently across periods. In case the state is high, both projects yield high outcomes, and if the state is low, low outcomes. In particular, the projects are (fully) correlated. Note, each project yields high and low outcomes with equal probabilities. Alternatively, consider another scenario in which the projects are independent (that is, the outcomes do not depend on the state of the economy) and both yield each outcome with equal probabilities. Since the agent needs to commit to exactly one project in each period, her investment strategy will depend on the marginal probabilities of each project. Thus, her investment strategy will be the same in these two scenarios. Putting differently, only the process of marginal distributions affect the optimal investment strategy; we need not worry about the effects of contemporaneous correlations on the optimal strategy.

This example is stylized; the specified projects are fully correlated and there is no learning (that is, the agent knows exactly what is the underlying distribution governing the state of the economy). We show, however, that the example can be generalized to any exploration problem where the outcome generating distribution (the object which the agent is attempting to learn) is stationary.² In particular, any such problem can be represented as another exploration problem in which the Bayesian model dictates that projects are not correlated, without affecting the agent's preferences over strategies, and hence the optimal strategy.

While this result allows an agent to simplify her decision problem, it also has a downside from the modeler's vantage. We show there is an inherent limitation on the type of beliefs that can be identified by observing an agent's preferences in such decision making environments; the general stochastic process governing beliefs can only be partially identified. Given an agent's preference over investment strategies in the two scenarios discussed above, there is no behavior that would identify which of the two scenarios

¹Exploration models were introduced by Robbins (1952) and have 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).

²We believe similar issues will arise for Markovian transitions. This seems to require different tools and we expect the characterizations to be different than in the stationary case.

the agent has in mind. Fortunately, even under the partial identification, we show there are meaningful behavioral restrictions allowing the modeler to test whether the agent is acting according to *some* Bayesian model.

Exploration Models and Elicited Beliefs. Because in exploration environments the agent can choose only one project in each period, her preferences over the different strategies depend only on the *margins* of her beliefs. And vice versa, the agent can only reveal—through choice or preference over investment strategies—her history dependent beliefs over each project separately. To show this, we provide a decision theoretic model, where we introduce a new dynamic and recursive framework capturing the exploration-exploitation tradeoffs faced by a decision maker (henceforth, DM). Our primitive is a preference relation over the different strategies that can be implemented by an agent facing a bandit problem. In this framework we first provide the behavioral (axiomatic) restrictions of subjective discounted expected utility maximization. The principal observation arising from our result is that the representation pins down only the processes of marginal beliefs of the different projects separately.

To better understand the economic relevance of this identification, we proceed with an analysis of a statistical framework. Here we consider stochastic processes that are determined by observing the outcome of a single project in each period, where potentially different projects are chosen across periods. We refer to these processes as *observable*, in light of the fact that they are precisely the output of our decision theoretic result. If, to the contrary, we had been able to observe the process over the joint realizations of all experiments, then the classic *exchangeability* property (or symmetry, as referred to in the decision theoretic terminology) would characterize Bayesianism. Given the limits of what can be observed, we cannot resort directly to such classical results. We provide instead a necessary and sufficient condition, *Across-Action Symmetry* (AA-SYM), for the observable processes to be consistent with an exchangeable process over the collection of all projects. It is not surprising, we obtain only a partial identification; a consistent exchangeable process, when it exists, need not be unique. We show however, that whenever our condition is met, there exists an exchangeable process, consistent with the observables, wherein there is no (contemporaneous) correlation across the different actions. Moreover, such a consistent exchangeable process *is* unique.

Finally, we show AA-SYM can be written in terms of the decision theoretic primitive. Combining these results, we conclude that the DM’s subjective joint distribution is not fully identified. Put differently, contemporaneous correlations across actions do not affect preferences and optimal strategies in bandit problems. Nonetheless, capitalizing on the sufficiency of AA-SYM for the marginals to be consistent with an underlying exchangeable model, we obtain an axiomatization for classic Bayesianism in *any* bandit problems.

The Decision Theoretic Model. In this framework a DM is tasked with ranking sequential and contingent choice objects: the action taken by the agent at any stage depends on the outcomes of previous actions. Formally, our primitive is a preference over *plans of action* (PoAs). Each action, a , is associated with a set of consumption prizes the action might yield, S_a . Then, a PoA is recursively defined as a *lottery over pairs* (a, f) , where a is an action and f is a mapping that specifies the continuation PoA for each possible outcome in S_a . Theorem 2 shows that the construction of PoAs is well defined. So, a PoA specifies an action to be taken each period that can depend on the outcome of all previously taken actions. See Figures 1 and 2, where $f(x), f(y), f(z)$ are themselves PoAs. Each node in a PoA can be identified by a *history* of

action-outcome realizations preceding it.

The actions in our model is in direct analogy to the arms of bandit problem (or actions in a repeated game). PoAs correspond to the set of all (possibly mixed) strategies in these environments. Note, however, the DM's perception of which outcome in S_a will result from taking action a is not specified. This is subjective and should be identified from the DM's preferences over PoAs. As discussed above, the main question is to what extent these beliefs can be identified and what are the economic implications of belief identification in this framework?




Figure 1: An action, a , and its support, S_a .




Figure 2: A degenerate PoA, (a, f) .

Theorem 3 axiomatizes preferences over PoAs of a DM who at each history entertains a belief regarding the outcome of future actions. That is, at each history h and for every action a , the DM entertains a belief $\mu_{h,a}$ over the possible outcomes S_a ; $\mu_{h,a}(x)$ is the DM's subjective probability that action a will yield outcome x , contingent on having observed the history h . Given this family of beliefs, the DM acts as a subjective discounted expected utility maximizer, valuing a PoA p , after observing h , according to a *Subjective Expected Experimentation (SEE)* representation:

$$U_h(p) = \mathbb{E}_p \left[\mathbb{E}_{\mu_{h,a}} [u(x) + \delta U_{h'}(f(x))] \right], \quad (\text{SEE})$$

where h' is the updated history (following h) when action a is taken and x is realized. All the parameters of the model –the consumption utility over outcomes, u , the discount factor, δ , and the history dependent subjective beliefs, $\{\mu_{h,a}\}_{h \in \mathcal{H}, a \in \mathcal{A}}$ – are identified uniquely.

Our setup requires a formulation of a novel axiom termed *proportionality (PRP)*: at any given history, the manner in which the DM evaluates continuation problems must be proportional to the manner in which she evaluates the consumption utility. Indeed, in order to ensure that the DM is acting consistently with a family of beliefs it must be that she assesses the value of each action according to the expectation of the consumption utility and discounted continuation utility it induces. Furthermore, it is necessary that the probabilistic weight she places on a given consumption utilities is the same as the weight she places on the corresponding continuation value.

Theorem 3 shows that PRP, along with (some of the) standard behavioral conditions for discounted expected utility, is necessary and sufficient for an SEE representation. While the axiomatization does not point to the optimal strategy in general strategic experimentation problems, which is known to be a hard problem to solve when actions are correlated, it provides (like most axiomatization theorems) a unifying guidance as to what might or might not be ruled out.

The identification result accompanying the representation concerns the marginal beliefs, $\{\mu_{h,a}\}_{h \in \mathcal{H}, a \in \mathcal{A}}$, and not a stochastic process over all actions, as is the starting point in the standard approach to bandit

problems. To be clear, this is not a limitation of the current setup any more than of bandit problems in general: observing a single pull of an arm each period simply does not provide sufficient data to identify the joint distribution. As such, we would like to know when the family of identified beliefs is consistent with an underlying exchangeable process. Further, given consistency, what are the limits of identification regarding this exchangeable model. To answer these questions we turn to the statistical model.

The Statistical Model. As above, there is a set of actions, \mathcal{A} , each element of which, a , is associated with the outcome space S_a . We are considering a family of processes over the outcomes of the different actions—where each period one and only one action is observed. Let $\mathbf{T} = (T_1, T_2, \dots)$, where $T_i \in \{S_a\}_{a \in \mathcal{A}}$ for every i . Let \mathcal{T} denote the set of all such sequences. For any \mathbf{T} in \mathcal{T} , let $\zeta_{\mathbf{T}}$ be a distribution over $\mathbf{T} = (T_1, T_2, \dots)$. We refer to these distributions as our observables; and, denoting $S = \prod_{a \in \mathcal{A}} S_a$, we assume a distribution, ζ , over $S^{\mathbb{N}}$ is not observable.

Our interest in this setup is motivated by the decision theoretic identification of $\{\mu_{h,a}\}_{h \in \mathcal{H}, a \in \mathcal{A}}$. While a process ζ over $S^{\mathbb{N}}$ specifies each $\mu_{h,a}$ (as the ζ -probability that following history h , action a will yield outcome x), it conveys strictly more information. For example, the ζ -probability that action a yields outcome x at the same time that action b yields outcome y has no counterpart in the identified family of marginals. The family of processes $\{\zeta_{\mathbf{T}}\}_{\mathbf{T} \in \mathcal{T}}$, on the other hand, contains exactly the same information as $\{\mu_{h,a}\}_{h \in \mathcal{H}, a \in \mathcal{A}}$. In both models, we do not have direct access to the probability of joint realizations of different actions. We only have access to the marginal distributions. Therefore, the exercise at hand concerns a direct translation of the decision theoretic observables into the statistical language.

In this framework, we introduce a condition referred to as *Across-Action Symmetry (AA-SYM)* and Theorem 4 shows that it is necessary and sufficient for the observables to be consistent with an exchangeable process over the joint realizations of all actions in every period (that is, $S^{\mathbb{N}}$). Informally, AA-SYM states that the probability of obtaining outcome x when taking action a followed by outcome y when taking action b , is the same as the probability of obtaining outcome y when taking action b followed by outcome x when taking action a . This is reminiscent of the symmetry (exchangeability) property, but note, in each period the outcome space may change as different actions can be taken.

The inherent observability constraint in this framework bears a cost; the exchangeable process with which our observables are consistent is typically not unique. Bearing in mind this generic non-uniqueness, we introduce what we term *strongly exchangeable* processes—a subclass of the widely studied exchangeable processes. We elaborate. Assume there is an underlying distribution governing the joint realization of actions that is *inter-temporally* i.i.d. This distribution is not known exactly, but there exists a prior probability over what it might be. The prior is updated every period upon the observation of the realization of actions. Due to [de Finetti \(1931\)](#); [Hewitt and Savage \(1955\)](#), these classical Bayesian updating processes are referred to as exchangeable. In a strongly exchangeable process, where the periodic state-space takes a product structure, the set of possible underlying distributions are such that outcomes across actions are independent. Thus, a strongly exchangeable process is one in which, conditional on the distributional parameter, outcomes are both inter-temporally and *contemporaneously* independent.

Despite strong exchangeability having more structure than classic exchangeability, it imparts no additional restrictions in our statistical model. Theorem 6 shows that a family of observables satisfying AA-SYM is consistent with a strongly exchangeable process, and this process is unique. We conclude, strong exchangeability is the full characterization of Bayesianism in our statistical framework and the lack of contemporaneous

correlations carry no constraints beyond AA-SYM.

Finally, returning to our decision theoretic model, we show that AA-SYM can be represented as an axiom over the primitives. Proposition 7 states that the additional axiom is both necessary and sufficient for the beliefs of an SEE representation to follow an exchangeable processes. This, of course, implies that two decision makers whose beliefs (as exchangeable processes) induce the same family of observable processes, will have the same preferences over strategies in *any* Bandit problem. In other words, Proposition 7 implies that contemporaneous correlations across actions do not impose any additional restrictions beyond classic Bayesianism when analyzing bandit problems, and no behavior can identify them in such an environment.

Organization. The paper is broadly broken into the two halves outlined above; Section 2 contains the decision theoretic framework, and Sections 3 and 4 the statistical one. Within Section 2, we first formally introduce the environment and the construction of plans of action (Sections 2.1). Next, we provide the axioms and representation result for an SEE structure (Sections 2.2 and 2.3). Section 3 introduces the observable processes that represent SEE belief structures. Here, we provide a statistical condition on observable processes, AA-SYM, so that the SEE belief structure is consistent with an exchangeable process. Section 4 introduces the notion of strong exchangeability and presents our (non) uniqueness result. The translation of AA-SYM back into decision theoretic terms is presented in Section 4.1. Section 5.1 discusses the related literature. An informal discussion regarding how a decision theoretic model would incorporate exogenous information appears in Section 5.2. Lastly, Section 5.3 discusses the point of disagreement among Bayesians in environments of experimentation. All the proofs are in the [Appendices](#).

2 THE DECISION THEORETIC FRAMEWORK

2.1 CHOICE OBJECTS

Constructing Plans of Action. The purpose of the current section is to construct the different choice objects, termed *plans of action (PoAs)*. The primitive of our model, as presented in the subsequent section, is a preference relation over all PoAs.

Let X be a finite set of outcomes, endowed with a metric d_X . Outcomes are consumption prizes. For any metric space, M , let $\mathcal{K}(M)$ denote the set non-empty compact subsets of M , endowed with the Hausdorff metric. Likewise, for any metric space M , denote $\Delta^{\mathcal{B}}(M)$ as the set of Borel probability distributions over M , endowed with the weak*-topology, and $\Delta(M)$ the subset of distributions with denumerable support.

Let \mathcal{A} be a compact and metrizable set of actions. Each action, a , is associated with a set of outcomes, $S_a \in \mathcal{K}(X)$, which is called the support of the action. We assume the map $a \mapsto S_a$ is continuous and surjective. For any metric space M , let $\mathcal{A} \otimes M = \{(a, f) | a \in \mathcal{A}, f: S_a \rightarrow M\} = \{(a, \{(x_i, m_i)\}_{i \in I}) \in \mathcal{A} \times \mathcal{K}(X \times M) | \bigcup_{i \in I} \{x_i\} = S_a \text{ and } x_i \neq x_j, \forall i \neq j \in I\}$, endowed with the subspace topology inherited from the product topology. By the continuity of $a \mapsto S_a$ we know that the relevant subspace is closed and hence the topology on $\mathcal{A} \otimes M$ is compact whenever M is. We can think of f as the assignment into M for each outcome in the support of action a . For any $f: X \rightarrow M$ we will abuse notation and write (a, f) rather than $(a, f|_{S_a})$.

With these definitions we can define PoAs. A PoA is a tree of actions, such that each period the DM receives a lottery (with denumerable support) over actions conditional on the outcomes for each of the previous actions.

We begin by constructing slightly more general objects. Set $R_0 = \Delta^{\mathcal{B}}(\mathcal{A})$; a 0-period plan is a lottery over actions. Given an action, an element of its support is realized and the plan is over. Then a 1-period plan, r_1 , is a lottery over actions, and continuation mappings into 0-period plans: $r_1 \in \Delta^{\mathcal{B}}(\mathcal{A} \otimes P_0)$. Given the realization of an action-continuation pair, (a, f) , in the support of r_1 , and the realized element of the support, $x \in S_a$, the DM receives a 0-period plan, as given by $f(x)$. Continuing in this fashion, we can define recursively,

$$R_n = \Delta^{\mathcal{B}}(\mathcal{A} \otimes R_{n-1}).$$

Define $R^* = \prod_{n \geq 0} R_n$. R^* is the set of all PoAs (including inconsistent plans and plans whose support is arbitrary). We first restrict ourselves to the set of *consistent* elements of R^* : those elements such that, the $(n-1)$ -period plan implied by the n -period plan is the same as the $(n-1)$ -period plan. To see why this is important, consider an element, r_n , of R_n . The plan r_n specifies an action to be taken in period 0 and, conditional on the outcome, the plan r_{n-1} which itself specifies the action to be taken in the next period and the continuation plan r_{n-2} for the next, etc. If we stop this process at any period $m < n$, ignoring whatever continuation plans are assigned, the output is an m period plan. Hence, each n period plan specifies a (unique) m period plan for each $m < n$. Moreover, an element $r \in R^*$ specifies an n period plan for each $n \in \mathbb{N}$. Intuitively, we would like to view each r as an infinite plan, by considering the sequence of arbitrarily large, and expanding, finite plans. Consistency is the requirement that makes this work, that for $r_n = \text{proj}_n r$, the first $m < n$ periods specify exactly $r_m = \text{proj}_m r$. Let R denote the set of all consistent plans.³

Proposition 1. *There exists a homeomorphism, $\lambda : R \rightarrow \Delta^{\mathcal{B}}(\mathcal{A} \otimes R)$ such that*

$$\text{marg}_{\mathcal{A} \times \mathcal{K}(X \times R_{n-1})}(\lambda(r)) = \text{proj}_n r. \quad (1)$$

Proof. In appendix A. ■

Next we want to consider plans whose support is denumerable. It is easy enough to set $P_0 = \Delta(\mathcal{A}) \subset R_0$, and define recursively $P_n = \Delta(\mathcal{A} \otimes P_{n-1}) \subset R_n$. Of course, there is a potential pitfall still lurking: for a given $\prod_{n \geq 0} P_n$, although each p_n is a denumerable lottery, the associated element, $\lambda(p)$ might live in $\Delta^{\mathcal{B}}(\mathcal{A} \otimes P)$ rather than $\Delta(\mathcal{A} \otimes P)$. Indeed, we need also to restrict our attention to the set of plans that have countable support not just for each finite level, but also “in the limit,” and whose implied continuation plans are also well behaved in such a manner. Fortunately, this can be done:⁴

Theorem 2. *There exists maximal set $P \subset R$ such that for each $p \in P$, $\text{proj}_n p \in P_n$, and λ is a homeomorphism between P and $\Delta(\mathcal{A} \otimes P)$.*

Proof. In appendix A. ■

³Precisely specifying the $m < n$ period plan implied by r_n requires a more cumbersome notation than we wish to introduce in the text; for the formal definitions see Appendix A.

⁴One can also consider measurable lotteries (instead of lotteries with countable support). In fact, the construction of the homeomorphism in Appendix A considers measurable lotteries. In the paper we focus on discrete support for notational cleanliness (see footnote 6) and tractability (to avoid measurability issues in proofs). We justify our focus by noting that $\Delta(\mathcal{A} \otimes P)$ is dense in $\Delta^{\mathcal{B}}(\mathcal{A} \otimes R)$ and so, given continuity (Axiom vNM), preferences over the more general objects are recoverable.

Notation	Meaning	Notes
$x, y, z \in X$	Single period consumption prizes	
$a \in A$	actions	associated with the support S_a
$(a, f) \in \mathcal{A} \otimes P$	action-continuation pair	$a \in A$ and $f : S_a \rightarrow P$
$p, q \in P$	PoA	Lottery over action-continuation pairs
$h \in \mathcal{H}(p)$	a history	corresponds to a unique node of p .

Figure 3: List of notational conventions.

The set P is our primitive.⁵ As a final notational comment, we would like to consider a further specification of *objective* plans, denoted by $\Sigma \subset P$. Σ denotes the set of plans which contain no subjective uncertainty; in every period, every possible action yields some outcome with certainty. Recall, for each $x \in X$ there is an associated action, a_x such that $S_{a_x} = \{x\}$. Associate this set of actions with X . Then $\Sigma_0 = \Delta(X)$ and, recursively, $\Sigma_n = \Delta(X \times \Sigma_{n-1})$. Finally $\Sigma = P \cap \prod_{n \geq 0} \Sigma_n$. That is, these plans specify only actions with deterministic outcomes at every stage. It is straightforward to show λ takes Σ to $\Delta(X \times \Sigma)$.

Historics. PoAs are infinite trees; each node, therefore, is itself the root of a new PoA—a distribution over action-continuation pairs. Each action-continuation, (a, f) , in the support of a node contains branches to new nodes (PoAs). The branches emanating from an action coincide with the outcomes in the support of that action, $x \in S_a$. The node that follows x is the PoA specified by $f(x)$. Each node, therefore, is reached after a unique history: the history specifies the realization of the distribution of each previous node, and outcome of the action realized. Thus, for a given PoA, p , each history of length n is an element of $\prod_{t=1}^n P \times [\mathcal{A} \otimes P] \times X$ such that $p^1 = p$ and

$$\begin{aligned} (a^t, f^t) &\in \text{supp}(p^t) \\ x^t &\in S_{a^t} \\ p^{t+1} &= f^t(x^t) \end{aligned}$$

Define the set of all histories of length n for p as $\mathcal{H}(p, n)$ and the set of all finite histories as $\mathcal{H}(p)$. Let $\mathcal{H}(n) = \bigcup_{p \in P} \mathcal{H}(p, n)$ and, $\mathcal{H} = \bigcup_{n \in \mathbb{N}} \mathcal{H}(n)$. For each $h \in \mathcal{H}(p, n)$, h corresponds to the node (PoA) defined by $f^n(x^n)$. Lastly, for any $p, q \in P$ and $h \in \mathcal{H}(p)$ define $p_{-h}q$ as the (unique!) element of P that coincides with p everywhere except after h in which case $f^n(x^n)$ is replaced by q . Note that the n period plan implied p and $p_{-h}q$ are the same. For any $p, q \in P$ and $n \in \mathbb{N}$, let $p_{-n}q \equiv \bigcup_{h \in \mathcal{H}(p, n)} p_{-h}q$.

Finally, for any $h = (p^1, a^1, f^1, x^1 \dots p^n, a^n, f^n, x^n)$ and $\hat{h} = (\hat{p}^1, \hat{a}^1, \hat{f}^1, \hat{x}^1 \dots \hat{p}^n, \hat{a}^n, \hat{f}^n, \hat{x}^n)$ both in $\mathcal{H}(n)$, we say that h and \hat{h} are \mathcal{A} -equivalent, denoted by $h \stackrel{\mathcal{A}}{\sim} \hat{h}$ if $a^i = \hat{a}^i$ and $x^i = \hat{x}^i$ for $i \leq n$. That is, two histories of length n are \mathcal{A} -equivalent, whenever they correspond to the same sequence of action-realization pairs, ignoring the objective randomization stage of each period and the continuation assignment

⁵One might consider an alternative framework of “adapted processes” of Anscombe-Aumann acts (see, for example, Epstein and Schneider (2003)), modified to our multi-action environment. In such a setup there is a distinction between exogenous states and outcomes (of the different actions). However, in a classical exploration problem, an outcome of an action is *simultaneously* an object from which the agent derives utility and from which the agent learns regarding the uncertainty underlying the (different) action(s). Similar results to those presented here would be obtained had we adopted the framework of adapted processes, but it seems conceptually appropriate to resort to plans of actions.

to outcomes that did not occur. It will turn out, we are only interested in the \mathcal{A} -equivalence classes of histories. Technically, this is the consequence of the linearity of preference and indifference to the resolution of uncertainty (as shown in Lemma 3); conceptually, this is because all uncertainty in the model regards the realization of actions, and so, observing objective lotteries has no informational benefit.

2.2 THE AXIOMS

The primitive in our model is a preference relation $\succsim \subseteq P \times P$ over all PoAs. When specific PoA and history are fixed, the preferences induce history dependent preferences as follows: for any $p \in P$, and $h \in \mathcal{H}(p)$ define $\succsim_h \subseteq P \times P$ by

$$q \succsim_h r \iff p_{-h}q \succsim p_{-h}r.$$

The following axioms will be employed over all history induced preferences.⁶ A history is *null* if \succsim_h is a trivial relation. This first four axioms are variants on the standard fare for discounted expected utility. They guarantee the expected utility structure, non-triviality, stationarity and separability (regarding objects over which learning cannot take place), respectively.

A1. (vNM). *The binary relation, \succsim_h satisfies the expected utility axioms. That is: weak order, continuity (defined over the relevant topology, see Appendix A) and independence.*

We require a stronger non-triviality condition that is standard, because of the subjective nature of the dynamic problem. We need to ensure the DM believes *some* outcome will obtain. Therefore, not all histories following a given action can be null.

A2. (NT). *For any non-null h , and any (a, f) , not all $h' \in h \times \mathcal{H}((a, f), n)$ are null.*

Of course, the nature of the problem at hand precludes stationarity and separability in full generality. Since the objective is to let the DM's beliefs depend on prior outcomes explicitly, her preferences will as well. However, the DM's beliefs do not influence her assessment of objective plans (i.e., elements of Σ), and so it is over this domain that stationarity and separability are retained. This means, the DM's preferences *in utility terms* are stationary and separable, but we still allow the conversion between actions and utils to depend on her beliefs which change responsively.

A3. (SST). *For all non-null $h \in \mathcal{H}$, and $\sigma, \sigma' \in \Sigma$,*

$$\sigma \succsim \sigma' \iff \sigma \succsim_h \sigma'.$$

A4. (SEP). *For all $x, x' \in X, \rho, \rho' \in \Sigma$ and $h \in \mathcal{H}$,*

$$\left(\frac{1}{2}(x, \rho) + \frac{1}{2}(x', \rho')\right) \sim_h \left(\frac{1}{2}(x, \rho') + \frac{1}{2}(x', \rho)\right).$$

⁶It is via the use of this construction that our appeal to denumerably supported lotteries provides tractability. If we were to employ lotteries with uncountable support, then histories would, in general, be zero probability events; under the expected utility hypothesis, \succsim_h would be null for all $h \in \mathcal{H}$. This could be remedied by appealing to histories as *events* in \mathcal{H} , measurable with respect to the filtration induced by previous resolutions of lottery-action-outcome tuples. We believe that this imposes a unnecessary notational burden.




Figure 4: A PoA, p , defined by $p(a, g) = \alpha$, $p(a, g') = \beta$ and $p(b, g'') = \gamma = 1 - \alpha - \beta$.




Figure 5: The PoA, $p.f$ where $f : X \rightarrow P$, and p is defined in figure 4. Notice, $p.f(a, f) = p(a, g) + p(a, g') = \alpha + \beta$.

Because of the two-stage nature of the resolution of uncertainty each period (first, the resolution of lottery over $\mathcal{A} \otimes P$, and then the resolution of the action over X), we need an additional separability constraint. From the point of view of period n , and when considering the continuation problem beginning in period $n + 1$, the DM should not care if uncertainty is resolved in period n (when the action-continuation pair is realized), or in period $n + 1$. That is, we also assume the DM is indifferent to the timing of objective lotteries *given a fixed action*.

A5. (IT). For all $a \in \mathcal{A}$, $h \in \mathcal{H}$, $\alpha \in (0, 1)$, and $(a, f), (a, g) \in \hat{P}$,

$$\alpha(a, f) + (1 - \alpha)(a, g) \sim_h (a, \alpha f + (1 - \alpha)g),$$

where mixtures of f and g are taken point-wise.

Thus far the axioms introduced are somewhat standard. However, in our particular framework these assumptions do not guarantee that the value of the action is in any way related with its realization of consumption alternatives. This is because, unlike other environments, the set of outcomes, X , plays a dual role in exploration models: representing both the space of outcomes and the state space regarding future actions.

The realization of an outcome x delivers utility according to both of these roles, and, to ensure consistency between them requires two steps. First, construct a subjective distribution over each action by treating X as a state space. This will be done by looking at the ranking of continuation mappings for each action (i.e., (a, f) compared to (a, g)). Interpreting X as the periodic state space, these continuation mappings are analogous to “acts” in the standard subjective expected utility paradigm—and so, standard techniques allow for the identification of such a subjective belief. Second, we need to ensure that the value assigned to arbitrary PoAs is the expectation according to these beliefs. Towards this, the following notation is introduced.

Definition. For any function $f : X \rightarrow \mathcal{P}$, define $p.f \in P$ as $p.f[(a, g)] = p[\{(b, h) | b = a\}]$ if $g = f$, and $p.f[(a, g)] = 0$ if $g \neq f$.

Take note, because we are dealing with distributions of denumerable support, we have no measurability concerns. The plan of action $p.f$ has the same distribution over actions in the first period, but the continuation plan is unambiguously assigned by f , as shown in Figures 4 and 5. If the original plan is in $\mathcal{A} \otimes P$,

then the dot operation is simply a switch of the continuation mapping: $(a, g).f = (a, f)$. This operation is introduced because it allows us to isolate the subjective distribution of the first period's action.

Definition. $p, q \in P$ are ***h -proportional*** if for all $f, g : X \rightarrow \Sigma$.

$$p.f \succsim_h p.g \iff q.f \succsim_h q.g$$

Since the images of f and g are in Σ , there is no informational effect from observing the outcome of p . Hence, f and g can be thought of as objective assignments into continuation utilities. The ranking ' $p.f \succsim p.g$ ' is really a ranking over f and g as functions from $X \rightarrow \mathbb{R}$. Thus, h -proportionality states that the DM's subjective uncertainty regarding X is the same when faced with p or with q .⁷

A6. (PRP). For all $p, q \in P$, and $f : X \rightarrow \Sigma$ if p and q are h -proportional then $p.f \sim_h q.f$.

The outcomes of an action represent not only the uncertainty regarding continuation, but also the utility outcome for the current period. So, when p and q are h -proportional, and thus induce the same uncertainty regarding X , the DM's uncertainty about her current period utility is the same across the plans. Therefore, if we replace the continuation problems with objectively equivalent plans, the DM should be indifferent between p and q .

2.3 A REPRESENTATION RESULT AND BELIEF ELICITATION

The following is our general axiomatization result. It states that the properties above characterize a DM who, when facing a PoA, calculates the subjective expected utility according to a collection of history dependent beliefs over action-outcome pairs, and among different PoAs contemplates the benefits of consumption versus learning.

Theorem 3 (Subjective Expected Experimentation Representation). \succsim_h satisfies vNM, NT, SST, SEP, IT and PRP if and only if there exists a utility index $u : X \rightarrow \mathbb{R}$, a discount factor $\delta \in (0, 1)$, and a family of beliefs $\{\mu_{h,a} \in \Delta(S_A)\}_{h \in \mathcal{H}, a \in A}$ such that

$$U_h(p) = \mathbb{E}_p \left[\mathbb{E}_{\mu_{h,a}} \left[u(x) + \delta U_{h'(a,x)}(f(x)) \right] \right], \quad (\text{SEE})$$

jointly represents $\{\succsim_h\}_{h \in \mathcal{H}}$, where $h'(a, x) = (h, p, (a, f), x)$. Moreover, u is cardinally unique, δ is unique, the family of beliefs is unique, and $\mu_{h,a} = \mu_{h',a}$ whenever $h \stackrel{A}{\sim} h'$.

Proof. In Appendix C. ■

The theorem states that we can (uniquely) elicit the beliefs, following every history, over the outcomes of each action separately. We will henceforth refer to such beliefs as an *SEE belief structure*. The axioms do not impose any restrictions on the dynamics of such beliefs. More importantly, the theorem shows that, when ranking the different strategies in a bandit problem, the decision maker does not reveal her beliefs over the *joint* realizations of the different actions.

⁷To see this, note that the relation R on $\mathbb{R}^X \times \mathbb{R}^X$ defined by fRg if and only if $p.f \succsim p.g$ is a preference relation over acts that satisfies the Anscombe and Aumann (1963) axioms, and therefore encodes the DM's subjective likelihood of each $E \subset X$. From a functional standpoint, h -proportionality states the subjective distribution over X induced by p is the same as that induced by q .

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 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 (alternatively, 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 $\mathbf{T} = (T_1, T_2, \dots)$ 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 T_1 , followed by T_2 , followed by T_3 , etc. For a given history of outcomes $h \in (T_1, T_2, \dots, T_n)$, we denote $h \in \mathbf{T}$ if $\mathbf{T} = (T_1, T_2, \dots, T_n, T_{n+1}, \dots)$. Lastly, for a sequence of experiments $\mathbf{T} = (T_1, \dots, T_n, T_{n+1}, \dots)$ and a permutation $\pi : n \rightarrow n$, denote $\pi\mathbf{T} = (T_{\pi(1)}, \dots, T_{\pi(n)}, T_{n+1}, \dots)$.

We consider a family of processes, $\{\zeta_{\mathbf{T}}\}_{\mathbf{T} \in \mathcal{T}}$, indexed by the possible sequences of experiments, \mathcal{T} . For a given family, we require that for sequences of experiments $\mathbf{T}, \mathbf{T}' \in \mathcal{T}$, if there is some history, $h \in \mathbf{T} \cap \mathbf{T}'$, 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. The set of all families of processes that meet this condition is in bijection to the set of all SEE belief structures,⁸ which is exactly the output of Theorem 3. Thus, in case there is no confusion, we will refer to a family of processes $\{\zeta_{\mathbf{T}}\}_{\mathbf{T} \in \mathcal{T}}$ as an SEE belief structure as well.

Exchangeable Processes and Consistency. Let $\mathcal{S}_{\mathcal{A}} \equiv \prod_{a \in \mathcal{A}} S_a$, and $\mathcal{S} \equiv \prod_{n \geq 0} \mathcal{S}_{\mathcal{A}}$. \mathcal{S} represents the grand state-space; a state, s , determines the realization of each action in each period –an entity which is unobservable to the modeler.

Definition. An SEE belief structure $\{\zeta_{\mathbf{T}}\}_{\mathbf{T} \in \mathcal{T}}$ is **consistent** with $\zeta \in \Delta^{\mathcal{B}}(\mathcal{S})$ if $\zeta|_{\mathbf{T}} = \zeta_{\mathbf{T}}$ for every $\mathbf{T} \in \mathcal{T}$.

That is, $\{\zeta_{\mathbf{T}}\}_{\mathbf{T} \in \mathcal{T}}$ is consistent with some process ζ over \mathcal{S} if for every sequence of experiments \mathbf{T} , the marginal of ζ to \mathbf{T} coincides with $\zeta_{\mathbf{T}}$. 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 bandit problems, we will pay particular attention to the class of *exchangeable* processes.

⁸Indeed, fix a family of history dependent beliefs $\{\mu_{h,a}\}_{h \in \mathcal{H}, a \in \mathcal{A}}$ and consider a sequence $\mathbf{T} = (S_{a_1}, S_{a_2}, \dots)$. Let $h = (x_1, x_2, \dots, x_n) \in \mathbf{T}$ and define

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

Then $\zeta_{\mathbf{T}}$ is defined as the unique (continuous) processes over \mathbf{T} that is consistent with $\zeta_{\mathbf{T}}(h)$ for every history $h \in \mathbf{T}$. This procedure can be inverted: Fix, $\{\zeta_{\mathbf{T}}\}_{\mathbf{T} \in \mathcal{T}}$ and some history $h = (a_1, x_1) \dots (a_n, x_n)$. Let \mathbf{T} be any sequence such that $T_i = S_{a_i}$ for $i \leq n$ and $T_{n+1} = S_a$. Then define

$$\mu_{h,a}(x) = \zeta_{\mathbf{T}}(T_{n+1} = x | T_1 = x_1 \dots T_n = x_n).$$

From here it is clear why we need to impose the condition that the probability of an event is not affected by future experiments: this condition arises naturally in the conditions of the SEE representation (see, Theorem 3). Recall that $\mu_{h,a} = \mu_{h',a}$ whenever $h \stackrel{\mathcal{A}}{\sim} h'$. Further, it is easy to check the above maps are continuous so that the bijection is in fact a homeomorphism.

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 there exists probability measure θ over $\Delta^{\mathcal{B}}(\mathcal{S}_{\mathcal{A}})$, such that

$$\zeta(E) = \int_{\Delta^{\mathcal{B}}(\Omega)} \hat{D}(E) d\theta(D), \quad (2)$$

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

Remark 1. If ζ is exchangeable, then θ is unique.

Exchangeable processes were first characterized by [de Finetti \(1931, 1937\)](#) and later extended by [Hewitt and Savage \(1955\)](#). Their fundamental result states that a process $\zeta \in \Delta^{\mathcal{B}}(\hat{\Omega})$ is exchangeable if and only if for any finite permutation $\pi : \mathbb{N} \rightarrow \mathbb{N}$ and event $E = \prod_{n \in \mathbb{N}} E_n$, we have

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

Exchangeable processes are of clear statistical importance, in particular within the subjectivist paradigm (see, for example [Schervish \(2012\)](#)). From the economic vantage, a DM who understands there to be an exchangeable process governing the outcome of actions would be considered Bayesian.⁹ This is because, given the representation in Eq. 2, the DM (acts as if she) entertains a second order distribution, which she updates following every observation.

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-Arm Symmetric (AA-SYM)** if

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

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

Intuitively, AA-SYM requires that if we consider a different order of experiments, then the probability of outcomes (in the appropriate order) does not change. **should we add the two coins examples here?** The next theorem states that across-arm symmetry 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 4. An SEE belief structure $\{\zeta_{\mathbf{T}}\}_{\mathbf{T} \in \mathcal{T}}$ satisfies AA-SYM if and only if it is consistent with an exchangeable process $\zeta \in \Delta^{\mathcal{B}}(\mathcal{S})$.

Theorem 4 is stated without proof. Necessity is trivial and sufficiency will be a straightforward application of Theorem 6.

⁹It is possible to consider more general Bayesian models than exchangeable processes. At least for the case of independent actions, for example, it is not hard to adapt a local consistency axiom as in [Lehrer and Teper \(2015\)](#) that will imply that beliefs follow a general martingale process.

4 STRONG EXCHANGEABILITY AND CONTEMPORANEOUS CORRELATIONS

Unfortunately, AA-SYM is not sufficient to obtain a unique exchangeable process consistent with an SEE belief structure. This lack of identification stems directly from the inability to observe the DM's beliefs regarding *contemporaneous* correlations. Consider two coins, a and b , which can both take values in $\{H, T\}$. Both coins are flipped each period. Consider the following two governing processes, which are i.i.d. across time periods. (1) the coins are perfectly correlated (with equal probability on HH and TT), or (2) the coins are identical and independent (and both have equal probability on H and T). Notice, the two cases induce the same marginal distributions over each coin *individually*. Thus, if the modeler has access only to the DM's marginal beliefs, the two processes are indistinguishable.

In this section we introduce a strengthening of exchangeability, which we aptly call *strongly exchangeable*, under which stochastic independence is preserved both inter-temporally (as in vanilla exchangeability) and *contemporaneously*.¹⁰

Definition. A process $\zeta \in \Delta^{\mathcal{B}}(\mathcal{S})$ is **strongly exchangeable** 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} .

Under a strongly exchangeable process 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 . Since exchangeable processes were first characterized as being invariant to permutations, for the sake of completeness we provide a similar characterization of strongly exchangeable processes.

Theorem 5. The process $\zeta \in \Delta^{\mathcal{B}}(\mathcal{S})$ is strongly exchangeable if and only 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

$$\zeta(E) = \zeta\left(\prod_{n \in \mathbb{N}} \prod_{a \in \mathcal{A}} E_{\pi_a(n), a}\right). \tag{4}$$

Proof. In Appendix C. ■

Following the intuition above, it should come as no surprise that under AA-SYM strong exchangeability can never be ruled out. In other words, there is no SEE belief structure, therefore no preferences over PoAs, that distinguishes exchangeability from strong exchangeability. Strongly exchangeable processes are ones where each dimension can be permuted independently. If $\pi_a = \pi_b$ for all $a, b \in \mathcal{A}$, the condition is exactly exchangeability. Strongly exchangeable process 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.

¹⁰We feel reasonably certain that strong exchangeability must have been studied previously in the statistics literature. However, we have found no references.

Theorem 6. *An SEE belief structure $\{\zeta_{\mathbf{T}}\}_{\mathbf{T} \in \mathcal{T}}$ satisfies AA-SYM if and only if it is consistent with a strongly exchangeable process. Furthermore, such a strongly exchangeable process is unique.*

Proof. In Appendix C. ■

4.1 AA-SYM AS A BEHAVIORAL RESTRICTION

In this section we introduce the axiomatic counterpart of AA-SYM, and so we can identify Bayesianism in exploration environments directly from preferences over the strategies.

Definition. *Let π be an n -permutation and $p, q \in P$. We say that q is π -permutation of p if for all $h \in \mathcal{H}(p, n)$, $h' \in \mathcal{H}(q, n)$, $\text{proj}_{\mathcal{A}^n} h = \pi(\text{proj}_{\mathcal{A}^n} h')$.*

If p admits any π -permutations it must be that the first n actions are assigned unambiguously (i.e., it does not depend on the realization of prior actions nor the objective randomization).

A7. (AA-SYM). *Let π be an n -permutation and $p, p' \in P$ with p' a π -permutation of p . Then, for all $a \in \mathcal{A}$, $\tau, \sigma, \sigma' \in \Sigma$, and $h \in \mathcal{H}(p, n)$, $h' \in \mathcal{H}(p', n)$, if h is a permutation of h' then*

$$p_{-n}\tau \succcurlyeq (p_{-n}\sigma)_{-h}\sigma' \iff p_{-n}\tau \succcurlyeq (p_{-n}\sigma)_{-h'}\sigma'.$$

After n periods the plan $p_{-n}\tau$ provides τ with certainty, while the plan $(p_{-n}\sigma)_{-h}\sigma'$ provides σ unless the history h occurs. Hence, the DM's preference between the plans depends on their ex-ante subjective assessment of how likely h is to occur. Similarly to the logic behind h -proportionality, AA-SYM states that the DM's assesses h to be exactly as probable as h' . In other words, the DM's likelihood of outcome realizations is invariant to the order in which the actions are taken. The intuition behind the next result is correspondingly straightforward.

Proposition 7 (Correlated Arms, Exchangeable Process). *Let \succcurlyeq admit an SEE representation with the associated observable processes $\{\zeta_{\mathbf{T}}\}_{\mathbf{T} \in \mathcal{T}}$. Then, the following are equivalent:*

1. \succcurlyeq_h satisfies AA-SYM;
2. $\{\zeta_{\mathbf{T}}\}_{\mathbf{T} \in \mathcal{T}}$ satisfies AA-SYM;
3. $\{\zeta_{\mathbf{T}}\}_{\mathbf{T} \in \mathcal{T}}$ is consistent with an exchangeable process; and
4. $\{\zeta_{\mathbf{T}}\}_{\mathbf{T} \in \mathcal{T}}$ is consistent with a (unique) strongly exchangeable process.

Proof. The proof that condition 1 is equivalent to condition 2 is provided Appendix C. Conditions 2, 3, and 4 are equivalent due to Theorem 6. ■

The proposition implies that strong-exchangeability carries no additional restrictions, beyond those of exchangeability, on agents' preferences over the different strategies in bandit problems, and in particular on their optimal strategies.

5 FURTHER DISCUSSION

5.1 RELATED LITERATURE

Within decision theory, the literature on learning broadly considers how a DM incorporates new information, generally via notions of Bayesianism and Exchangeability, and often in the domain of uncertainty: see [Epstein and Le Breton \(1993\)](#); [Epstein and Seo \(2010\)](#); [Klibanoff et al. \(2013\)](#); [Lehrer and Teper \(2015\)](#). Recently, there has been an interest in subjective learning, or, the identification of the set of possible “signals” that the DM believes she might observe. At it’s most simple, this is the elicitation of the set of potential tastes (often referred to as subjective states) the decision maker anticipates, accomplished by examining the DM’s preference over *menus* of choice objects: see [Kreps \(1979\)](#); [Dekel et al. \(2001\)](#). By also incorporating consumption goods that contract on an objective state space, the modeler can interpret the DM’s preference for flexibility as directly stemming from her anticipation of acquiring information regarding the likelihood of states, as in [Dillenberger et al. \(2014\)](#); [Krishna and Sadowski \(2014\)](#).

There is also a small but highly relevant literature working on the identification of responsive learning. [Hyogo \(2007\)](#) considers a two-period model, with an objective state space, in which the DM ranks action-menu pairs. The action is taken in the first period and provides information regarding the likelihood of states, after the revelation of which, the DM choose a state-contingent act from the menu. The identification of interest is the DM’s subjective interpretation of actions as signals. Similarly, [Cooke \(2016\)](#) entertains a similar model without the need for an objective state-space, and in which the consumption of a single object in the first period plays the role of a fully informative action. Cooke, therefore, identifies both the state-space and the corresponding signal structure. [Piermont et al. \(2016, forthcoming\)](#) consider a recursive and infinite horizon version of Kreps’ model, where the DM deterministically learns about her preference regarding objects she has previously consumed. [Dillenberger et al. \(2015\)](#) consider a different infinite horizon model where the DM makes separate choices in each period regarding her information structure and current period consumption. It is worth pointing out, all of these models, unlike the this paper, capitalize on the “preference for flexibility” paradigm to characterize learning. We are able to identify subjective learning without appealing to the menu structure because of the purely responsive aspect of our model. In other words, flexibility is “built in” to our setup, as a different action can be taken after every possible realization of the signal (action).

5.2 SUBJECTIVE LEARNING WITH ENDOGENOUS AND EXOGENOUS INFORMATION

As witnessed the literature covered above, there seems to be a divide in the literature regarding subjective learning. In one camp, are models that elicit the DM’s perception of exogenous flows of information (as a canonical example, take [Dillenberger et al. \(2014\)](#)), and in the other are models that assume information is acquired only via actions taken by the DM (where this paper lies). Realistically, neither of these information structures capture the full gamut of information transmission in economic environments.

Consider the following example within the setup of the current paper. A firm is choosing between two projects (actions), a and b . Assume that each project has a high-type and a low type. The firm believes (after observing h) the probability that each project is the high-type is $\mu_{h,a}$ and $\mu_{h,b}$, respectively. By experimenting between a and b the firm’s beliefs and preferences will evolve.

But, what happens if the firm anticipates the release of a comprehensive report regarding project a just before period 1? This report will declare project a high quality with probability $\alpha^h > \frac{1}{2}$ if the projects true

type is high and with probability $\alpha^l < \frac{1}{2}$ if it is low. Hence, the report is an informative signal. Now, if the firm's belief after observing h in period 0 is given by $[\mu_{h,a}, \mu_{h,b}]$ then, according to Bayes rule, the firm's belief regarding project a being the high-type, at the beginning of period 1 will be $\mu_{h,a}^+ = \frac{\alpha^h \cdot \mu_{h,a}}{\alpha^h \cdot \mu_{h,a} + \alpha^l (1 - \mu_{h,a})}$, if the report is positive, and $\mu_{h,a}^- = \frac{(1 - \alpha^h) \cdot \mu_{h,a}}{(1 - \alpha^h) \cdot \mu_{h,a} + (1 - \alpha^l) \cdot (1 - \mu_{h,a})}$ if the report is negative.

Unfortunately, however, the ex-ante elicitation of preferences in our domain cannot capture the anticipation of information. The firm is ranking PoAs according to its aggregated belief from the ex-ante perspective, and thus, so as to maximize its expected belief:

$$(\alpha^h \mu_{h,a} + \alpha^l (1 - \mu_{h,a})) \mu_{h,a}^+ + ((1 - \alpha^h) \mu_{h,a} + (1 - \alpha^l) (1 - \mu_{h,a})) \mu_{h,a}^- = \mu_{h,a}.$$

Because of the Bayesian structure, the DM's beliefs must form a martingale, so her expectation of her anticipated beliefs are exactly her ex-ante beliefs. This fact, coupled with the linearity of expected utility, imply that the DM's ex-ante preference over PoAs is unaffected by her anticipation of exogenous information arrival.

All hope is not lost, however, of fully characterizing the DM's subjective information structure. The approach of [Dillenberger et al.](#) is orthogonal to our's, leading us to conjecture that the two models can co-exist and impart a clean separation between exogenous and endogenous information flows. Going back to the example, imagine there are two PoAs, p and q such that p is preferred to q under beliefs μ_h^+ , and q to p under μ_h^- . The DM would therefore strictly desire flexibility after period 0, even after she is able to condition her decision on h . Of course, because the report is released after period 0, irrespective of the action taken by the DM, for any 0-period history h' , there must exist some other PoAs, p' and q' , for which flexibility is strictly beneficial (after h').

5.3 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 6 states that **AA-SYM** is necessary and sufficient for an **SEE** belief system to be consistent with some exchangeable process. As discussed in the Introduction, **AA-SYM** projected to stochastic processes is weaker than the standard symmetry axiom applied in the literature, because the standard assumption requires that histories fully specify the evolution of the state, while in our setup, histories can only specify cylinders. Because **AA-SYM** is a weaker assumption, de Finetti's theorem implies that processes satisfying such an assumption need not be exchangeable and have a Bayesian representation as in Eq. (2).

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 **AA-SYM**, but the process itself is clearly not exchangeable. Nevertheless, Theorem 6 guarantees that there is a (unique) strongly exchangeable process that is consistent with the **SEE** belief structure. In this case it is easy to see that that process would be the one in which every period we

toss two coins that are identical and independent (and both have equal probability on H and T).

A ON THE CONSTRUCTION OF PLANS.

Generalized Plans. We will begin by constructing a more general notion of Plans (reminiscent of IHCPs, first constructed in [Gul and Pesendorfer \(2004\)](#)), and then refine our notion to capture only the elements of interest. This methodology serves two purposes. First, the more general approach allows us to use standard techniques for the construction of infinite horizon choice objects. Second, generalized plans may be of direct interest in future work, when, for example, denumerable support is not desirable. To begin, let $Q_0 = \Delta^{\mathcal{B}}(\mathcal{A})$ and, for define recursively for each $n \geq 1$

$$Q_n = \Delta^{\mathcal{B}}(\mathcal{A} \times \mathcal{K}(X \times Q_{n-1})).$$

Finally, define $Q^* = \prod_{n \geq 0} Q_n$. Q^* is the set of generalized plans.

Consistency. For the first step, we will follow closely GP of consistent IHCPs, but with enough difference that it makes sense to define things outright. Formally, let $G_1 : \mathcal{A} \times \mathcal{K}(X \times Q_0) \rightarrow \mathcal{A}$ as the mapping $(a, \{x, q_0\}) \mapsto a$. Let $F_1 : Q_1 \rightarrow Q_0$ as the mapping $F_1 : q_1 \mapsto (E \mapsto q_1(G_1^{-1}(E)))$, for any $E \in \mathcal{B}(\mathcal{A})$. Therefore, for any $E \in \mathcal{B}(\mathcal{A})$, $F_1(p_1)(E)$ is the probability of event E in period 0 as implied by p_1 ; $F_1(p_1)$ is the distribution over period 0 actions implied by p_1 . From here we can recursively define $G_n : \mathcal{A} \times \mathcal{K}(X \times Q_n) \rightarrow \mathcal{A} \times \mathcal{K}(X \times Q_{n-1})$ as:

$$G_n : (a, \{x, q_{n-1}\}) \mapsto (a, \{x, F_{n-1}(q_0)\})$$

and $F_n : Q_n \rightarrow Q_{n-1}$ as:

$$F_n : q_n \mapsto (E \mapsto q_n(G_n^{-1}(E)))$$

for any E in $\Delta^{\mathcal{B}}(\mathcal{A} \times \mathcal{K}(X \times Q_{n-1}))$. A consistent generalized plan is one such that

$$F_n(q_n) = q_{n-1},$$

for all n . Let Q denote all such generalized plans.

A.1 CONSTRUCTION PROOFS

Lemma 1. *There exists a homeomorphism, $\lambda : Q \rightarrow \Delta^{\mathcal{B}}(\mathcal{A} \times \mathcal{K}(X \times Q))$ such that*

$$\text{marg}_{\mathcal{A} \times \mathcal{K}(X \times Q_{n-1})}(\lambda(q)) = q_n. \tag{5}$$

Proof. [STEP 1: EXTENSION THEOREM.] Let $C_n = \{(q_0, \dots, q_n) \in \prod_{k=0}^n Q_k \mid q_k = F_{k+1}(q_{k+1}), \forall k = 1 \dots n-1\}$, and $T_n = \mathcal{K}(X \times C_n)$ for $n \geq 0$. Let $T^* = \prod_{n=0}^{\infty} T_n$ and $T = \{t \in T^* \mid (\text{proj}_{T_n} t_{n+1} = t_n)\}$. Let $Y_0 = \Delta^{\mathcal{B}}(\mathcal{A})$ and for $n \geq 1$ let $Y_n = \Delta^{\mathcal{B}}(\mathcal{A} \times T_0 \times \dots \times T_n)$. We say the the sequence of probability measures $\{\nu_n \in Y_n\}_{n \geq 0}$ is consistent if $\text{marg}_{\mathcal{A} \dots T_{n-1}} \nu_{n+1} = \nu_n$ for all $n \geq 0$. Let Y^c denote the set of all consistent sequences. Then we know by [Brandenburger and Dekel \(1993\)](#), for every $\{\nu_n\} \in Y^c$ there exists a unique $\nu \in \Delta^{\mathcal{B}}(\mathcal{A} \times T^*)$ such that $\text{marg}_{\mathcal{A}} \nu = \nu_0$ and $\text{marg}_{\mathcal{A} \dots T_n} \nu = \nu_n$. Moreover, the map $\psi : Y^c \rightarrow \Delta^{\mathcal{B}}(\mathcal{A} \times T^*)$:

$$\psi : \{\nu_n\} \mapsto \nu$$

is a homeomorphism. \square

[STEP 2: EXTENDING BACKWARDS.] Let $D_n = \{(t_0, \dots, t_n) \in \times_{k=0}^n T_n \mid t_k = \text{proj}_{T_n}(t_{k+1}), \forall k = 1 \dots n-1\}$. Let $Y^d = \{\{\nu_n\} \in Y^c \mid \nu_n(\mathcal{A} \times D_n) = 1, \forall n \geq 0\}$. We will now show, for each $q \in Q$, there exists a unique $\{\nu_n\} \in Y^d$, such that $\nu_0 = q_0$ and $\text{marg}_{\mathcal{A} \times \mathcal{K}(X \times Q_{n-1})}(\text{marg}_{\mathcal{A} \times T_{n-1}}(\nu_n)) = q_n$ for all $n \geq 1$. Indeed, let m_0, m_1 be the identify function on \mathcal{A} and $\mathcal{A} \times \mathcal{K}(X \times Q_0)$, respectively. Then for each $n \geq 2$ let $m_n : \mathcal{A} \times D_{n-1} \rightarrow \mathcal{A} \times \mathcal{K}(X \times Q_{n-1})$ as follows:

$$m_{n+1} : (a, \{x^0, q_0^0\}, \{x^1, q_0^1, q_1^1\} \dots \{x^n, q_0^n \dots q_n^n\}) \mapsto (a, \{x^n, q_n^n\}).$$

Note: for $n \geq 0$, each m_n is a Borel isomorphism. Indeed, continuity of m_n is obvious, and measurability follows immediately from the fact that canonical projections are measurable in the product σ -algebra. It is clear that m_n is surjective, and —since (given F_k for $k \in 1 \dots n$) q_n uniquely determines $q_0 \dots q_{n-1}$, which, (given the projection mappings) uniquely determines $T_0 \dots T_{n-1}$ — m_i is also injective. As for, m_n^{-1} , continuity follows from the continuity of F_k for $k \in 1 \dots n$ and the projection mappings. Lastly, measurability of m_n^{-1} comes from the fact that a continuous injective Borel function is a Borel isomorphism (see [Kechris \(2012\)](#) corollary 15.2).

So, let $\psi : Q \rightarrow Y^d$ as the map

$$\phi : q \mapsto \{E_n \mapsto q_n(m_n(E_n))\}_{n \geq 0},$$

for any $E_n \in \mathcal{B}(\mathcal{A} \times T_0 \times \dots \times T_n)$. The continuity of ϕ and ϕ^{-1} follow from the fact that they are constructed from the pushforward measures of m_n^{-1} and m_n , respectively, which are themselves continuous (or, explicitly, see GP lemma 4).

Finally, let $\Gamma_n = \mathcal{A} \times D_n \times \times_{k=n+1}^\infty T_k$. Let $\nu = \psi(\{\nu_n\})$ for some $\{\nu_n\}$ in Y^d . Then $\nu(\Gamma_n) = \nu(\mathcal{A} \times D_n) = 1$. So, $\nu(\mathcal{A} \times T) = \nu(\cap_{n \geq 0} \Gamma_n) = \lim \nu(\Gamma_n) = 1$. Also, note, if $\nu(\mathcal{A} \times T) = 1$, then $\nu(\Gamma_n) = 1$ for all $n \geq 0$. So, $\nu \in Y^d$ if and only if $\nu(\mathcal{A} \times T) = 1$, i.e., if, $\psi(Y^d) = \{\nu \in \Delta^{\mathcal{B}}(\mathcal{A} \times T^*) \mid \nu(\mathcal{A} \times T) = 1\}$. \square

[STEP 3: EXTENDING FORWARDS.] Let τ denote the map from $\mathcal{A} \times \mathcal{K}(X \times Q) \rightarrow \mathcal{A} \times T$ as

$$\tau : (a, \{x, q\}) \mapsto (a, (\{x, q_0\}, \{x, q_0, q_1\}, \dots))$$

That τ is a bijection follows from the consistency conditions on Q, T , and C_n for $n \geq 1$. Now takes some measurable set $E \subseteq T$. Then $\tau^{-1}(E) = \bigcap_{n \geq 0} \{\{x, q_0, \dots, q_n \times_{k=n+1}^\infty Q_k\} \in \mathcal{K}(X \times Q^*)\}$, the countable intersection of measurable sets, and hence measurable. That τ and τ^{-1} are continuous is immediate. Therefore, by the same argument as in [STEP 2], τ is a Borel isomorphism and $\kappa : \Delta^{\mathcal{B}}(\mathcal{A} \times T) \rightarrow \Delta^{\mathcal{B}}(\mathcal{A} \times \mathcal{K}(X \times Q))$,

$$\kappa : \nu \mapsto (E \mapsto \nu(\tau(E)))$$

for all E in $\Delta^{\mathcal{B}}(\mathcal{A} \times \mathcal{K}(X \times Q))$. Clearly, $\text{marg}_{\mathcal{A}}(\kappa(\nu)) = \text{marg}_{\mathcal{A}}(\nu)$ and $\text{marg}_{\mathcal{A} \times \mathcal{K}(X \times Q_{n-1})}(\kappa(\nu)) = \text{marg}_{\mathcal{A} \times \mathcal{K}(X \times Q_{n-1})}(\nu)$ for all $n \geq 1$. \square

Behold, $\lambda = \kappa \circ \psi \circ \phi$ is the desired homeomorphism. \blacksquare

Definition. Let $R_0 = Q_0$ and $R_1 = \{r_1 \in Q_1 \mid r_1(\mathcal{A} \otimes R_0) = 1\}$. Then, recursively let $R_n = \{r_n \in Q_n \mid r_n(\mathcal{A} \otimes R_{n-1}) = 1\}$. Set $R = \prod_{n=0}^\infty R_n$.

Proof of Proposition 1. We show that λ is a homeomorphism between R and $\Delta^{\mathcal{B}}(\mathcal{A} \otimes R)$. Identify $\Delta^{\mathcal{B}}(\mathcal{A} \otimes R)$ with $\{\nu \in \Delta^{\mathcal{B}}(\mathcal{A} \times \mathcal{K}(X \times Q)) \mid \nu(\mathcal{A} \otimes R) = 1\}$. Let $r \in R$. For each $n \geq 0$ let $\Gamma_n^r = \{(a, \{x, q\}) \in$

$\mathcal{A} \otimes Q | q_k \in R_k, k = 0 \dots n \}$. Then $\lambda(r)(\Gamma_n^r) = \text{marg}_{A \times K(X \times Q_n)}(\lambda(r))(\mathcal{A} \otimes R_n) = r_{n+1}(\mathcal{A} \otimes R_n) = 1$ for all $n \geq 1$. So $\lambda(r)(\mathcal{A} \otimes R) = \lambda(r)(\cap_{n \geq 0} \Gamma_n^r) = \lim \lambda(r)(\Gamma_n^r) = 1$. Now, fix $q \in Q$ with $\lambda(q)(\mathcal{A} \otimes R) = 1$, then $q_n(\mathcal{A} \otimes R_{n-1}) = \text{marg}_{A \times K(X \times Q_{n-1})}(\lambda(q))(\mathcal{A} \otimes R_{n-1}) = \lambda(r)(\Gamma_n^r) \geq \lambda(r)(\mathcal{A} \otimes R) = 1$ for all $n \geq 0$ and so $q \in R$. \blacksquare

Definition. For a metric space, M , let $\Delta(M) \subseteq \Delta^B(M)$ denote the set of all distributions with countable support. I.e., for all $\nu \in \Delta(M)$, there exists a countable set S_ν such that $m \notin S_\nu \implies \nu(m) = 0$, and $\sum_{m \in S_\nu} \nu(m) = 1$.

Definition. Set $W : \mathcal{P}(R) \rightarrow \mathcal{P}(R)$ as the function:

$$W : E \mapsto \{r' \in R | r' \in \text{Im}(f) \text{ for some } (a, f) \in \text{supp}(\lambda(r)), r \in E\}$$

Definition. Let $P_0 = \Delta(\mathcal{A})$ and $P_1 = \{p_1 \in R_1 | p_1 \in \Delta(\mathcal{A} \otimes P_0)\}$. Then, recursively let $P_n = \{p_n \in R_n | p_n \in \Delta(\mathcal{A} \otimes P_{n-1})\}$. Set $P = \{p \in \prod_{n=0}^{\infty} P_n | \prod_{n=0}^{\infty} \lambda(W^n(r)) \subset \prod_{n=0}^{\infty} \Delta(\mathcal{A} \otimes R)\}$.

Proof of Theorem 2. We show that λ is a homeomorphism between P and $\Delta(\mathcal{A} \otimes P)$. First note, by construction, for all $r \in R$, $\lambda(r) \in \Delta^B(\mathcal{A} \otimes W(r))$. Let $p \in P$; by the conditions on P , $\lambda(p) \in \Delta(\mathcal{A} \otimes R)$. Therefore, it suffices to show that for any $p \in P$, and $r \in W(p)$, $r \in P$. So fix some $r \in W(p)$. It follows from an analogous argument to Corollary 1 that $r \in \prod_{n=0}^{\infty} P_n$. Finally, note that $W^{n-1}(r) \subseteq W^n(r)$, for all $n \geq 2$. \blacksquare

B LEMMAS.

Lemma 2. If \geq_h satisfies vNM and IT, then \geq_h satisfies the sure thing principal:

A8. (STP). For all $a \in \mathcal{A}$ and $f, f', g, g' : X \rightarrow P$, such that, for all $x \in X$, either (i) $f(x) = f'(x)$ and $g(x) = g'(x)$ or (ii) $f(x) = g(x)$ and $f'(x) = g'(x)$. Then,

$$(a, f) \geq_h (a, g) \iff (a, f') \geq_h (a, g').$$

Proof. Assume this was not true and, without loss of generality, that $(a, f) \geq_h (a, g)$ but $(a, g') >_h (a, f')$. Now notice, when mixtures are taken point-wise, $\frac{1}{2}f + \frac{1}{2}g' = \frac{1}{2}g + \frac{1}{2}f'$. Therefore,

$$\begin{aligned} \left(\frac{1}{2}(a, f) + \frac{1}{2}(a, g')\right) &>_h \left(\frac{1}{2}(a, g) + \frac{1}{2}(a, f')\right) \\ &\sim_h \left(a, \frac{1}{2}g + \frac{1}{2}f'\right) = \left(a, \frac{1}{2}f + \frac{1}{2}g'\right) \\ &\sim_h \left(\frac{1}{2}(a, f) + \frac{1}{2}(a, g')\right), \end{aligned}$$

where the first line follows from vNM, and the indifference conditions from IT. This is a contradiction. \blacksquare

Lemma 3. If \geq_h satisfies vNM and IT for all $h \in \mathcal{H}$, then, if $h \stackrel{A}{\sim} h'$ then $\geq_h = \geq_{h'}$.

Proof. We will show the claim on induction by the length of the history. So let $h, h' \in \mathcal{H}(1)$ such that $h \stackrel{A}{\sim} h'$. Therefore, $h = (p, (a, f), x)$ and $h' = (p', (a, g), x)$. Notice, by definition we have, $p = \alpha(a, f) + (1 - \alpha)r$ and $p' = \alpha'(a, g) + (1 - \alpha')r'$, for some $\alpha, \alpha' \in (0, 1]$ and $r, r' \in P$.

Let $q, q' \in P$; we want to show that $q \succcurlyeq_h q' \iff q \succcurlyeq_{h'} q'$. So let $q \succcurlyeq_h q'$, or by definition, $p_{-h}q \succcurlyeq p_{-h}q'$, which by the above observation is equivalent to

$$\alpha(a, f)_{-((a, f), (a, f), x)}q + (1 - \alpha)r \succcurlyeq \alpha(a, f)_{-((a, f), (a, f), x)}q + (1 - \alpha)r.$$

By independence (i.e., **vNM**) this is if and only if $(a, f)_{-((a, f), (a, f), x)}q \succcurlyeq (a, f)_{-((a, f), (a, f), x)}q'$, which by **STP** is if and only if $(a, g)_{-((a, g), (a, g), x)}q \succcurlyeq (a, g)_{-((a, g), (a, g), x)}q'$. Using independence again, this is if and only if $p'_{-h'}q \succcurlyeq p'_{-h'}q'$. This completes the base case.

So assume the claim holds for all histories of length n . So let $h, h' \in \mathcal{H}(n+1)$ such that $h \stackrel{A}{\sim} h'$. Therefore, $h = (h_n, p, (a, f), x)$ and $h' = (h'_n, p', (a, g), x)$, for some $h_n, h'_n \in H(n)$ such that $h_n \stackrel{A}{\sim} h'_n$. By the inductive hypothesis $\succcurlyeq_{h_n} = \succcurlyeq_{h'_n}$.

Let $q, q' \in P$, and $q \succcurlyeq_h q'$, or by definition, $p_{-(p, (a, f), x)}q \succcurlyeq_{h_n} p_{-(p, (a, f), x)}q'$. By independence and the sure thing principle this is if and only if $(a, g)_{-((a, g), (a, g), x)}q \succcurlyeq_{h_n} (a, g)_{-((a, g), (a, g), x)}q'$, which by independence again (and the equivalence of \succcurlyeq_{h_n} and $\succcurlyeq_{h'_n}$), is if and only if $p'_{-(p', (a, g), x)}q \succcurlyeq_{h'_n} p'_{-(p', (a, g), x)}q'$. ■

C PROOF OF MAIN THEOREMS

Proof of Theorem 3. [**STEP 0: VALUE FUNCTION.**] Since \succcurlyeq_h satisfies **vNM**, there exists a $v_h : \mathcal{A} \otimes P \rightarrow \mathbb{R}$ such that

$$U_h(p) = \mathbb{E}_p [v_h(a, f)] \tag{6}$$

represents \succcurlyeq_h , with v_h unique up to affine translations. □

[**STEP 1: RECURSIVE STRUCTURE.**] To obtain the skeleton of the representation, let's consider $\hat{\succcurlyeq}$, the restriction of \succcurlyeq to Σ (i.e., using the natural association between streams of lotteries and degenerate trees). The relation $\hat{\succcurlyeq}$ satisfies **vNM** (it is continuous by the closure of Σ in P). Hence there is a linear and continuous representation: i.e., an index $\hat{u} : X \times \Sigma \rightarrow \mathbb{R}$ such that:

$$\hat{U}(\sigma) = \mathbb{E}_\sigma [\hat{u}(x, \rho)] \tag{7}$$

unique up to affine translations.

Following [Gul and Pesendorfer \(2004\)](#), (henceforth **GP**), fix some $(x', \rho') \in \Sigma$. From **SEP** we have $\hat{U}(\frac{1}{2}(x, \rho) + \frac{1}{2}(x', \rho')) = \hat{U}(\frac{1}{2}(x, \rho') + \frac{1}{2}(x', \rho))$, and hence, $\hat{u}(x, \rho) = \hat{u}(x, \rho') + \hat{u}(x', \rho) - \hat{u}(x', \rho')$. Then setting $u(x) = \hat{u}(x, \rho') - \hat{u}(x', \rho')$ and $W(\rho) = \hat{u}(x', \rho)$, we have,

$$\hat{U}(\sigma) = \mathbb{E}_\sigma [u(x) + W(\rho)] \tag{8}$$

Now, consider $p' = (x', \rho)$. Notice that p' has unique 1-period history: $h = (p', p', x')$. By **NT**, h cannot be null. So, by **SST**, $\hat{\succcurlyeq}_h = \hat{\succcurlyeq}$. This implies, of course that $W = \delta \hat{U} + \beta$ for some $\delta > 0$ and $\beta \in \mathbb{R}$. Following **Step 3** of [Lemma 9](#) in **GP** exactly, we see that $\delta < 1$ and without loss of generality we can set $\beta = 0$:

$$\hat{U}(\sigma) = \mathbb{E}_\sigma [u(x) + \delta \hat{U}(\rho)] \tag{9}$$

Both representing $\hat{\succcurlyeq}$ and being unique up to affine translations, we can normalize each U_h to coincide with \hat{U} over Σ . □

[**STEP 2: THE EXISTENCE OF SUBJECTIVE PROBABILITIES.**] For each $a \in \mathcal{A}$ consider

$$\mathcal{F}(a) = a \otimes \Sigma$$

i.e., the elements of \hat{P} that begin with action a and from period 2 onwards are in Σ . Associate $\mathcal{F}(a)$ with the set of “acts”: $f : S_a \rightarrow \Sigma$, in the natural way. For any acts f, g let $f_{-x}g$ denote the act that coincides with f for all $x' \in S_a$, $x' \neq x$, and coincides with g after x . For each $h \in \mathcal{H}$, and acts $f, g \in \mathcal{F}(a)$, say $f \succsim_{h,a} g$ if and only if $(a, f) \succsim_h (a, g)$.

It is immediate that $\succsim_{h,a}$ is a continuous weak order (where, as before, continuity follows from the closure of \mathcal{F} in P). Further, $\succsim_{h,a}$ satisfies independence. Indeed: fix $f, g, h \in \mathcal{F}(a)$ with $f \succsim_{h,a} g$. Then

$$\begin{aligned} f \succsim_{h,a} g &\implies (a, f) \succsim_h (a, g) \\ &\implies \alpha(a, f) + (1 - \alpha)(a, h) \succsim_h \alpha(a, g) + (1 - \alpha)(a, h) \\ &\implies (a, \alpha f + (1 - \alpha)h) \succsim_h (a, \alpha g + (1 - \alpha)h) \\ &\implies \alpha f + (1 - \alpha)h \succsim_{h,a} \alpha g + (1 - \alpha)h, \end{aligned}$$

where the third line uses IT. Lastly, $\succsim_{h,a}$ satisfies monotonicity, a direct consequence of SST and STP. Hence, we have state-independence which gives us the full set of [Anscombe and Aumann \(1963\)](#) axioms for an SEU representation of $\succsim_{h,a}$ with state space S_a . That is, a belief $\mu_{h,a} \in \Delta(S_a)$ and a utility index from $\Sigma \rightarrow \mathbb{R}$ (which is of course, \hat{U} , and so will be denoted as such), such that

$$\hat{V}_{h,a}(f) = \mathbb{E}_{\mu_{h,a}} [\hat{U}(f(x))] \tag{10}$$

represents $\succsim_{h,a}$. □

[STEP 3: PROPORTIONAL ACTIONS.] Now, fix some $h \in \mathcal{H}$ and consider an arbitrary $(a, f) \in \mathcal{A} \otimes P$. Let $\rho \in \Sigma$ be such that $\text{marg}_X \rho = \mu_{h,a}$. We claim, (a, f) and ρ are h -proportional. Fix some $g, g' : X \rightarrow \Sigma$. From (10), we know

$$(a, g) \succsim_h (a, g') \iff \mathbb{E}_{\mu_{h,a}} [\hat{U}(g(x))] \geq \mathbb{E}_{\mu_{h,a}} [\hat{U}(g'(x))] \tag{11}$$

From (9) we have

$$\begin{aligned} \hat{U}(\rho.g) &= \mathbb{E}_\rho [u(x) + \delta \hat{U}(g(x))] \\ &= \mathbb{E}_{\text{marg}_X \rho} [u(x) + \delta \hat{U}(g(x))] \\ &= \mathbb{E}_{\mu_{h,a}} [u(x)] + \delta \mathbb{E}_{\mu_{h,a}} [\hat{U}(g(x))] \end{aligned}$$

In corresponding fashion we obtain the analogous representation for $\hat{U}(\rho.g')$, and hence

$$\rho.g \succsim_h \rho.g' \iff \mathbb{E}_{\mu_{h,a}} [\hat{U}(g(x))] \geq \mathbb{E}_{\mu_{h,a}} [\hat{U}(g'(x))] \tag{12}$$

Combining the implications of (11) and (12), we see that (a, f) and ρ are h -proportional. □

[STEP 4: PROPORTIONAL PLANS.] We now claim that for any $h \in \mathcal{H}$ and $p \in P$ there exists some $\sigma \in \Sigma$ such that $p \sim_h \sigma$. Fix some $p \in P$, and for each $n \in \mathbb{N}$ define p^n to be any PoA that agrees with p on the first n periods, then provides elements of Σ unambiguously. Note that $p_n \rightarrow p$ point-wise and hence converges in the product topology. Therefore, the claim reduces to finding a convergent sequence $\{\sigma_n\}_{n \in \mathbb{N}} \subset \Sigma$ such that $\sigma^n \sim_h p^n$, as continuity ensures the limits are indifferent.

We will prove the subsidiary claim by induction. Consider p^1 , for each $(a, f) \in \text{supp}[p^1]$, note, by assumption, $f : X \rightarrow \Sigma$. Let $\tau^{1,(a,f)} \in \Sigma$ be such that $\text{marg}_X \tau^{1,(a,f)} = \mu_{h,a}$. By [STEP 3], (a, f) and $\tau^{1,(a,f)}$ are h -proportional. And thus, $\tau^{1,(a,f)}.f \sim_h (a, f).f = (a, f)$, by PRP. Let $\sigma^1 \in \Sigma$ be such that

$\sigma^1[E] = p^1[\{(a, f)|\tau^{1,(a,f)}.f \in E\}]$. Therefore,

$$\begin{aligned} U_h(p^1) &= \mathbb{E}_{p^1} [v_h(a, f)] \\ &= \mathbb{E}_{p^1} [\hat{U}(\tau^{1,(a,f)}.f)] \\ &= \mathbb{E}_{\sigma^1} [\hat{U}(\rho)] \\ &= \hat{U}(\sigma^1) \end{aligned}$$

where the third line comes from the change of variables formula for pushforward measures. This completes the base case.

Now, assume the claim hold for all h and $m \leq n - 1$ for some $n \in \mathbb{N}$. Consider p^n . Note that for all h' of the form $h(x) = (h, p^n, (a, f), x)$, the implied continuation problem $p^n(h')$ satisfies the inductive hypothesis. Therefore, there exists a $\sigma^{n-1, h'} \sim_{h'} p(h')$ for all such h' .

Let \star denote the mapping: $(a, f) \mapsto (a, f)^\star = (a, x \mapsto \sigma^{n-1, h(a,x)})$, where $h(a, x) = (h, p^n, (a, f), x)$. By construction, for each (a, f) in $\text{supp}(p^n)$, and $x \in S_a$ we have $(a, f) \sim_h (a, f_{-x} \sigma^{n-1, h(a,x)})$ (using the notation from [STEP 2]). Employing STP we have $(a, f) \sim_h (a, f)^\star$ (i.e., enumerating the outcomes in S_a and changing f one entry at a time, where STP ensures that each iteration is indifferent to the last).

Let $\hat{p}^n \in P$ be such that $\hat{p}^n[E] = p^n[\{(a, f)|(a, f)^\star \in E\}]$. So,

$$\begin{aligned} U_h(p^n) &= \mathbb{E}_{p^n} [v_h(a, f)] \\ &= \mathbb{E}_{p^n} [v_h((a, f)^\star)] \\ &= \mathbb{E}_{\hat{p}^n} [v_h(b, g)] \\ &= U_h(\hat{p}^n) \end{aligned}$$

Applying the base case to \hat{p}^n concludes the inductive step. Notice also, the convergence of $\{\sigma^n\}_{n \in \mathbb{N}}$ is easily verified, following the fact that the marginals on p_n are fixed for any σ^m with $m \geq n$. \square

[STEP 5: REPRESENTATION.] Consider any $(a, f) \in \mathcal{A} \otimes P$. We claim that there exists an $(a, f') \in \mathcal{F}(a)$ such that $(a, f) \sim_h (a, f')$. Indeed, by [STEP 4], for any $x \in S_a$, there exists some $\rho(a, x)$ such that $\rho(a, x) \sim_{h(a,x)} f(x)$, where $h(a, x) = (h, (a, f), (a, f), x)$. Define $f' \in \mathcal{F}(a)$ as $x \mapsto \rho(a, x)$. It follows from STP that $(a, f) \sim_h (a, f')$.

We know by [STEP 3] that there exists a $\rho \in \Sigma$, h -proportional to (a, f) , with $\text{marg}_X \rho = \mu_{h,a}$. Hence $(a, g) = (a, f).g \sim_h \rho.g$ for all $g : X \rightarrow \Sigma$. We have,

$$\begin{aligned} v_h(a, g) &= \hat{U}(\rho.g) \\ &= \mathbb{E}_{\mu_{h,a}} [u(x) + \delta \hat{U}(g(x))], \end{aligned}$$

and so, for (a, f') :

$$v_h(a, f') = \mathbb{E}_{\mu_{h,a}} [u(x) + \delta \hat{U}(\rho(a, x))].$$

By the indifference condition $\rho(a, x) \sim_{h(a,x)} f(x)$,

$$v_h(a, f) = \mathbb{E}_{\mu_{h,a}} [u(x) + \delta U_{h(a,x)}(f(x))]. \quad (13)$$

Notice, $h(a, x) \stackrel{A}{\approx} h'(a, x) = (h, p, (a, f), x)$, so by Lemma 3, $\succsim_{h(a,x)} = \succsim_{h'(a,x)}$. Applying this fact, and plugging (13) into (6) provides

$$U_h(p) = \mathbb{E}_p [\mathbb{E}_{\mu_{h,a}} [u(x) + \delta U_{h'(a,x)}(f(x))]] \quad (14)$$

as desired. □

Proof of Theorem 5. First we show, if a strongly exchangeable process ζ over \mathcal{S} is induced by an i.i.d distribution D over $\mathcal{S}_{\mathcal{A}}$, 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 $E, F \in S_{\hat{\mathcal{A}}}$, $E', F' \in S_{\hat{\mathcal{A}}'}$, be measurable events. Identify E^n with the cylinder it E generates in \mathcal{S} when in the n^{th} coordinate: $E^n = \{s \in \mathcal{S} | s_{n, \mathcal{B}} \in E\}$. Since ζ is strongly exchangeable we have that

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

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 $S_{\mathcal{A}}$)

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

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 $E \in S_{\hat{\mathcal{A}}}$ and $E', F' \in S_{\hat{\mathcal{A}}'}$, 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 de Finetti's theorem and represent the process ζ by

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

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 $\zeta(\cdot | s_t) : S_{\mathcal{A}} \rightarrow [0, 1]$ 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 $\zeta(\cdot, \cdot | s_t) : S_{\mathcal{A}} \times S_{\mathcal{A}} \rightarrow [0, 1]$, 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. ■

Proof of Theorem 6. Fix an SEE belief structure $\{\hat{\zeta}_{\mathbf{T}}\}_{\mathbf{T} \in \mathcal{T}}$. We first construct a pre-measure $\hat{\zeta}$ on the semi-algebra of cylinder sets. Fix any 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 $n \in m_1 \dots m_1 + m_2$, $\hat{E}_{n,a} = S_a$ for all a except a_2 , etc. Let $\mathbf{T}(E)$ denote the sequence such that $T_n = S_i$ 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 $n \in m_1 \dots m_1 + m_2$, $T_2 = S_{a_2}$, 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 $\hat{\zeta}(E) = \zeta_{\mathbf{T}(E)}(\hat{E})$.

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 $E_{n,a}^k \subseteq E_{n,a}^{k+1}$. But each $E_{n,a}^k$ is finite, hence compact, and nonempty, because $\zeta(E^k) \geq \epsilon$. Therefore $\bigcap_{k \in \mathbb{N}} E_{n,a}^k \neq \emptyset$. 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 $\bar{\pi}_a$ denote a finite permutation for each $a \in \mathcal{A}$. Let $F = \prod_{n \in \mathbb{N}} \prod_{a \in \mathcal{A}} E_{\bar{\pi}_a(n), a}$. Let π_{a_i} denote the permutation given by the construction of \hat{F} . Then $\hat{F} = \prod_{n \in \mathbb{N}} \prod_{a \in \mathcal{A}} E_{\pi_a(\bar{\pi}_a(n)), a}$. In particular, 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 AA-SYM, $\zeta_{\mathbf{T}(\hat{E})}(\hat{E}) = \zeta_{\pi^* \mathbf{T}(\hat{E})}(\pi^* \hat{E}) = \zeta_{\mathbf{T}(\hat{F})}(\hat{F})$. Therefore, $\zeta(E) = \zeta(F)$ and so, by Theorem 5, 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 7. Let $\{\mu_{h,a}\}_{h \in \mathcal{H}, a \in \mathcal{A}}$ be an SEE structure for \geq that satisfies AA-SYM. Let $\{\zeta_{\mathbf{T}}\}_{\mathbf{T} \in \mathcal{T}}$ be the associated family of observable processes. Fix \mathbf{T} and some n period history $h \in \mathbf{T}$. Let, $(a_1, x_1) \dots (a_n, x_n)$, where for each $i \leq n$ let a_i is such that $T_i = S_{a_i}$ and x_i is the i^{th} component of h . This represents an \mathcal{A} -equivalence class of decision theoretic histories. In our standard abuse of notation, let h also denote this class of histories. Following this abuse, when it is not confusing to do so, let πh denote both the permuted statistical history and the \mathcal{A} -equivalence class represented by $(a_{\pi(1)}, x_{\pi(1)}) \dots (a_{\pi(n)}, x_{\pi(n)})$.

Fix some n -permutation π . Let p denote the PoA that assigns a_i in the i^{th} period with certainty. Let p' be the π -permutation of p . We have

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

Let $\sigma, \sigma' \in \Sigma$ be such that $U_h(\sigma) = 1$ and $U_h(\sigma') = 0$. Then, by (SEE) we have

$$p_{-n}(\alpha\sigma + (1 - \alpha)\sigma') \sim (p_{-n}\sigma')_{-h}\sigma$$

so, by AA-SYM, we have,

$$p'_{-n}(\alpha\sigma + (1 - \alpha)\sigma') \sim (p'_{-n}\sigma')_{-h'}\sigma$$

which implies, again by (SEE),

$$\alpha = \mu_{\emptyset, a_{\pi(1)}}(x_{\pi(1)}) \cdot \mu_{(a_{\pi(1)}, x_{\pi(1)}), a_{\pi(2)}}(x_{\pi(2)}) \cdots \mu_{(a_{\pi(1)}, x_{\pi(1)}, \dots, a_{\pi(n-1)}, x_{\pi(n-1)}), a_{\pi(n)}}(x_{\pi(n)}) = \zeta_{\pi\mathbf{T}}(\pi h).$$

Hence, $\zeta_{\mathbf{T}}(h) = \zeta_{\pi\mathbf{T}}(\pi h)$ as desired. ■

REFERENCES

- Francis J Anscombe and Robert J Aumann. A definition of subjective probability. *The annals of mathematical statistics*, 34(1):199–205, 1963.
- Dirk Bergemann and Juuso Välimäki. Bandit Problems. In Steven N. Durlauf and Lawrence E. Blume, editors, *The New Palgrave Dictionary of Economics*. Second edition, 2008.
- Donald A Berry and Bert Fristedt. *Bandit problems: sequential allocation of experiments (Monographs on statistics and applied probability)*. Springer, 1985.
- Adam Brandenburger and Eddie Dekel. Hierarchies of Beliefs and Common Knowledge. *Journal of Economic Theory*, 59(1):189–198, 1993.
- Kevin Cooke. Preference Discovery and Experimentation. 2016. *Working paper*.
- Bruno de Finetti. Funzione caratteristica di un fenomeno aleatorio. 1931.
- Bruno de Finetti. La prévision: ses lois logiques, ses sources subjectives. In *Annales de l'institut Henri Poincaré*, volume 7, pages 1–68, 1937.
- Eddie Dekel, Barton L. Lipman, and Aldo Rustichini. Representing Preferences with a Unique Subjective State Space. *Econometrica*, 69(4):891–934, 2001.
- David Dillenberger, Juan Sebastián Lleras, Philipp Sadowski, and Norio Takeoka. A theory of subjective learning. *Journal of Economic Theory*, pages 1–31, 2014.
- David Dillenberger, R. Vijay Krishna, and Philipp Sadowski. Dynamic Rational Inattention. 2015. *Working paper*.
- Larry G Epstein and Martin Schneider. Recursive multiple-priors. *Journal of Economic Theory*, 113(1): 1–31, 2003.
- Larry G Epstein and Kyoungwon Seo. Symmetry of evidence without evidence of symmetry. *Theoretical Economics*, 5(3):313–368, 2010.
- L.G. Epstein and M. Le Breton. Dynamically consistent beliefs must be Bayesian. *Journal of Economic Theory*, 61:1–22, 1993.

- Faruk Gul and Wolfgang Pesendorfer. Self-Control and the Theory of Consumption. *Econometrica*, 72(1):119–158, 2004.
- Edwin Hewitt and Leonard J. Savage. Symmetric measures on Cartesian products. *Transactions of the American Mathematical Society*, 80(2):470, 1955.
- Kazuya Hyogo. A subjective model of experimentation. *Journal of Economic Theory*, 133(1):316–330, 2007.
- Alexander Kechris. *Classical descriptive set theory*, volume 156. Springer Science & Business Media, 2012.
- Peter Klibanoff, Sujoy Mukerji, and Kyoungwon Seo. Perceived Ambiguity and Relevant Measures. *Econometrica*, 82(5):1–46, 2013.
- David M Kreps. A Representation Theorem for “Preference for Flexibility”. *Econometrica*, 47(3):565–577, 1979.
- R. Vijay Krishna and Philipp Sadowski. Dynamic Preference for Flexibility. *Econometrica*, 82(133):655–703, 2014.
- Ehud Lehrer and Roei Teper. Who is a Bayesian? 2015. *Working paper*.
- Evan Piermont, Norio Takeoka, and Roei Teper. *Games and Economic Behavior*, 2016, forthcoming.
- Herbert Robbins. Some aspects of the sequential design of experiments. *Bulletin of the American Mathematical Society*, 58(5):527–536, 1952.
- Mark J Schervish. *Theory of statistics*. Springer Science & Business Media, 2012.