Research in Economics (1999) **53**, 227–242
Article No. reec.1999.0186
Available online at http://www.idealibrary.com on IDEAL®

Bayesianism without learning

DOV SAMET†

Department of Economics, University College London, Gower Street, London WCIE 6DP, U.K.

Summary

According to the standard definition, a Bayesian agent is one who forms his posterior belief by conditioning his prior belief on what he has learned, that is, on facts of which he has become certain. Here it is shown that Bayesianism can be described without assuming that the agent acquires any certain information; an agent is Bayesian if his prior, when conditioned on his posterior belief, agrees with the latter. This condition is shown to characterize Bayesian models.

J.E.L. Classification: C72, D80, D83.
Keywords: Bayesianism, type spaces, learning.

1. Introduction

This paper studies the dynamics of beliefs, that is, the way agents change their prior belief into posterior belief. The standard assumption made in economics and game theory is that agents are Bayesian, which means that posterior beliefs are formed from prior beliefs by conditioning on acquired knowledge. The most commonly used models for Bayesian agents are Harsanyi type spaces. It is shown here that Bayesianism, as expressed by type spaces, can be interpreted differently. A change of belief, according to the suggested interpretation, is not necessarily a result of acquiring new information or knowledge. The only knowledge that is necessarily involved is knowledge of the change itself. We show that the description of Bayesian agents by type spaces can be characterized without assuming the acquisition of any knowledge leading to the change.

† The author acknowledges helpful comments by Pierpaolo Battigalli, Giacomo Bonanno, Joseph Halpern, Aviad Heifetz, Ehud Lehrer, Barry O'Neill, and Daniel Osherson.

1090–9443/99/020227 + 16 $30.00/0

After deriving these results, I discovered that similar results have already been obtained by Gaifman (1988). Here we put these results in the context of the economic literature on type spaces, and stress their implications to the nature of Bayesian change of belief. We discuss Gaifman's work, as well as other related works in Subsection 1.5.

1.1. BAYESIANISM DEFINED BY LEARNING

The Bayesian paradigm deals with the relation between an agent's present (*posterior*) and former (*prior*) belief. The common interpretation of Bayesianism, in a nutshell, is that beliefs change by learning: the agent comes to know, or at least becomes certain of, some facts. He, then, forms his posterior belief by conditioning his prior belief on these facts. There are several expressions, or metaphors, that are used to describe the agent's learning of facts. In the economics and decision theory literature an agent is said to observe a *signal*; in models of games with incomplete information a player is said to learn his *type*.

A formal model that depicts the relationship between the prior and the posterior belief of a Bayesian agent consists of a state space with a probability distribution on it which describes prior belief. Learning is introduced into the model by adding a certain partition to the state space. The posterior belief in a given element of the partition is the prior conditioned on this element. According to the common interpretation of Bayesianism, an element of the partition consists of all the states in which the agent observed a certain signal; each such element is an event that formalizes the notion "all the information that the agent has acquired".

The underlying assumption of this interpretation of Bayesianism—that beliefs change as a result of becoming certain of some facts—has been questioned and criticized by many students of subjective probability. Experience tells us, according to the critics' argument, that sometimes we change our beliefs without being able to specify any relevant facts, of which we became certain. Thus, for example, we may have some prior belief concerning the honesty of a person and change it considerably after a short conversation with him, in which no new facts are revealed to us. The change is the result of an immediate impression that cannot be reduced to or explained by the learning of any new fact about the person in discussion.

It seems that the very definition of a Bayesian agent, and the formal model of type space, cannot describe any change of belief that is not the result of gaining certainty concerning some facts. This model *presumes* the existence of events—the elements of the partition—of which the agent is certain. How can conditioning be carried out without such events? What would be the meaning of Bayesianism without conditioning one's belief on what has been learned?

This paper suggests that Bayesianism can be interpreted in a way that does not require that beliefs change as a result of learning. This interpretation is made possible by an alternative definition of a Bayesian agent, one which does not assume that the agent has learned and become certain about any facts or events. We show that the structure of the Bayesian model, as described above, is *implied* by our definition of a Bayesian agent, rather than being *assumed*. In particular, this new definition implies that our agent must be certain of some facts—he must be certain of his own belief. The elements of the partition can now be interpreted as describing the certainty the agent has about his own posterior belief. Thus, certainty under this interpretation is not the *cause* of the change of belief, but its *result*.

1.2. BAYESIANISM DEFINED WITHOUT LEARNING

We discuss now, informally, our proposed characterization of a Bayesian agent. We start by showing how such a characterization can be derived from the standard formulation of Bayesianism, which assumes that posterior beliefs are derived by conditioning prior beliefs on learned information. The latter can be roughly expressed by,

> posterior belief = prior belief as modified by the information
> the agent has learned.

But the information the agent has learned is just what led him to revise his prior belief, by conditioning, to the posterior one. Thus we can write,

> posterior belief = prior belief as modified by the information
> that made the agent change his prior to his posterior belief.

Now, the case that the agent has learned the information that makes him change his prior belief to some specified posterior belief is precisely the case that he holds this posterior belief. Indeed, the information leads to the posterior belief, but also conversely; this belief includes what the agent knows, or is certain of, that is, the information he has learned.

Thus, the event over which we condition is just the event that the agent holds his present posterior belief, and we can reformulate the relationship between prior and posterior beliefs as,

> posterior belief = prior belief as modified by the posterior belief.

We derived this latter relationship between prior and posterior from the standard definition of Bayesianism. But note that the formulation of this relationship does not assume learning, certainty, or knowledge of any event. It describes a certain kind of consistency that posterior and prior beliefs should satisfy, but not the way by which the posterior is derived from the prior. In this paper we show that this latter relationship can be used to characterize, or equivalently, to define Bayesianism. We prove that it implies the standard model of Bayesian agents.

1.3. AN ILLUSTRATION

To illustrate the proposed alternative definition, consider the following question we pose to John.

> "Given that *next month* you will believe the probability of Clinton completing a full term to be at least 0·3 and the probability of Saddam Hussein continuing to reign until the year 2010 to be at least 0·8, do you agree, *now*, that the probabilities of these events are at least 0·3 and 0·8, respectively?"

We ask John to evaluate his prior concerning two events *given* some bounds on his posterior belief concerning the same events. What we expect John to answer, according to the definition of Bayesianism proposed here, is just: "Yes." That is, the conditioned prior should conform with the given posteriors.

The reasoning that would lead John to give us this answer is simple. By conditioning on the belief he may have in a month, John puts himself, right now, in the same position he is conditioning on. What other answer could we expect?

Our characterization of Bayesianism is a straightforward formalization of the above dialogue with John; an agent is Bayesian if he answers in the affirmative to all questions of this form, as John does.†

1.4. THE MODEL

Formalizing the previous illustration requires a model in which we can identify the event that the posterior of a given event E is greater than or equal to some number p. This, in turn, is made possible by making the posterior depend on, and vary with, the points of the model. This is precisely the main feature of type spaces as defined by Harsanyi (1967–68); the type of an agent, that is, his

† The question posed to John concerns two events. As we show in Section 4, we can equivalently require that he answers in the affirmative when asked about any number of events. But it is not enough to ask him about only one event at a time.

(posterior) probability over the state space, depends on the state. In Harsanyi's type spaces, or other similar Bayesian models, it is assumed that the agent knows, or is certain of, his type. Here we do not want to make any assumption of this kind, because we want to *derive* the structure of Bayesian models, rather than *assume* it. Thus our models, which we call *belief spaces*, are simpler and more general. A belief space is a measurable space each point (or state) of which is associated with a probability measure on the space; no further restrictions other than simple measurability conditions are imposed.

1.5. RELATED LITERATURE

An important element of the principle introduced in the previous two sections is a conditional probability $p(E|C)$ where C describes the probability of some events for a probability measure q. Several authors studied principles that involve such conditioning. These principles, like the one we state here, require that the conditional probability be consistent with the probabilities described by C. Such principles are sometimes called Miller's principle, although Miller (1966), who first introduced a principle of this kind, considered it paradoxical.†

The various principles in this category differ in the interpretation of the probabilities p and q, and the specification of the condition C.‡ Here, p and q are the prior and posterior probabilities of an agent. This is also the interpretation given to these probabilities by van Fraassen (1984) who calls this principle Reflection. Lewis (1980) studied what he calls the Principal Principle, in which p is subjective probability, while q is objective probability. In Skyrms (1980), Halpern (1991, 1998), and Samet (1997, 1998*b*), the probabilities p and q are the same. Gaifman (1988) labels p as the probability of an agent, and q as that of an expert.

In this work, as well as in Gaifman (1988), C describes the probabilities of two (or more) events, one of which is E. In some works it is a full description of the probability q. In Skyrms (1980),

† Miller claimed that the principle, now bearing his name, contradicts the calculus of probability. Popper (1966), in a comment that follows Miller's short note, hailed the "important and indeed brilliant discovery described by Miller," and concluded that it entails the *"complete abandonment of the philosophical idea that there is an 'inductive logic'* which is, formally, identical with the calculus of relative probabilities—that is, with the *probabilistic generalization of deductive logic."* When the principle is stated in a rigorous mathematical model, the paradox seems to vanish.

‡ It is possible to state similar principles for non-quantitative beliefs. See, for example, Battigalli and Bonanno (1997) who study the condition that believing an event is equivalent to believing that it will be believed at a later time.

Halpern (1991, 1998), and Samet (1997, 1998*b*), where p and q coincide, C describes only the probability of E.

The way C describes the probability of an event also varies. Here it is given as a lower bound on the probability. In Gaifman (1988) and Halpern (1991, 1998) it is given by lower and upper bounds, and in Skyrms (1980) C describes the exact probability. These differences are technical and of minor importance.

Until Gaifman (1988), none of the works that discussed Miller's principle in its various forms proposed any set-theoretic model in which such principles can be studied rigorously. The main problem facing such modelling is the identification of C, which describes certain probabilistic statements, with an event. Harsanyi (1967–68) faced the same problem when he analysed games with incomplete information. The difficulty there stemmed from the need to describe beliefs of players about other players' beliefs. Hence, such beliefs have to be described by an event. His model, the type space, is defined as the product of type sets, one for each player, where a type of a player is defined as a probability distribution on the type of other players (or even the whole space with the restriction that the player is certain about his type). This model has been modified and simplified by Mertens and Zamir (1985) under the name belief space. Here we adopt a somewhat less restrictive model than theirs under the same name. In his modelling, Gaifman, unaware of the solution proposed by Harsanyi, arrived at the same solution, a model in which beliefs, given as probability distributions over a probability space, or using another terminology, a state space, vary with the states.

Although very similar results were obtained previously by Gaifman (1988), it is worthwhile to present this work for the following reasons. First, it is appropriate to restate these results in the context and terminology of type spaces and belief spaces; a model which has become so common in economics and game theory over the last three decades, and is the main tool for describing Bayesianism in these fields. Philosophers too could benefit greatly by using these models. Although Gaifman introduced his set-theoretic model in the philosophy literature, only a few papers in this area have used it.

Second, it is important to state the implications of these results for the dynamics of belief change in general and for learning in Bayesian theory in particular. Students of the dynamics of belief change have not been aware of these implications. For example, Maher (1993) dedicates a whole chapter in his book to belief change, but he does not make any reference to Gaifman (1988). Contrary to the result of Gaifman's work, he concludes that "It is possible for the shift from p to q to satisfy Reflection without it being the case that there is a proposition E such that $q(\cdot) = p(\cdot \mid E)$." In light of the interpretation we suggest here, a conclusion like this needs at least to be restricted appropriately.

1.6. SYNOPSIS

In the next section we formally introduce belief spaces. The first two subsections of Section 3 discuss the two features that make belief spaces Bayesian: that the agent knows his type, and that he forms his posterior belief by conditioning his prior on his type. Propositions 1 and 2 suggest new ways to express the well-known notions of type and of knowing-the-type, in terms of belief-describing events. The usefulness of such expressions is demonstrated in the proof of the main theorems. In theorems 1 and 2 in Section 4 we present two variants of the characterization of Bayesianism without learning. The proofs are given in Section 5.

2. Belief spaces

The model we use to formalize prior and posterior beliefs is a probability space—the probability being the prior—in which each point is associated with a probability over the space which is the posterior at that point. The association of points, or states, with probabilities is the basic idea that underlies Harsanyi *type spaces*. The details, as well as the name *belief spaces*, are more in tune with the account of type spaces in Mertens and Zamil (1985). Our model is simpler than theirs and than most type spaces in the literature, as no topology on the type space or the space of probability measures is required, in the spirit of the model in Heifetz and Samet (1998).

DEFINITION 1: *a* type space *is a quadruple* (Ω, Σ, μ, t) *where,*

(1) Ω *is a measurable space with a* σ*-field* Σ*, generated by a countable field* Σ_0*. The members of* Ω *are called* states*, and the members of* Σ *are called* events.
(2) μ *is a* σ*-additive probability measure on* Ω*, called the* prior.
(3) t *is a map from* Ω *to* $\Delta(\Omega)$ *—the set of all* σ*-additive probability measures on* Ω*—such that for each* $E \in \Sigma$*,* $t(\cdot)(E)$ *is a measurable real function. For a state* $\omega, t(\omega)$ *is called the* type*, or the* posterior *of the agent at* ω.

A special role is played here by events that describe the agent's belief. For each event E and real number p,

$$B^p(E) = \{\omega \mid t(\omega)(E) \geq p\}$$

is the event that the *probability ascribed by the agent to E is at least p*, or in short, the event that the *probability of E is at least p*. Events of the form $B^1(E)$ are important and deserve a special name. We call $B^1(E)$ the event that the agent is *certain* of E.

The measurability condition on t is tantamount to saying that for each E and p, $B^p(E)$ is a measurable set. When a type space, as well as $\Delta(\Omega)$, are endowed with a topology, it is the standard requirement that the function t be continuous that guarantees the measurability requirement on t.

Belief spaces as defined here are very general objects. No restrictions are imposed on the way beliefs are associated with different states. No requirements are imposed on what the agent should know, or be certain of. No relation between the prior and posterior beliefs is specified. Further structure is required in order to turn belief spaces into the Bayesian models used in economics and game theory. We adopt a general definition of belief spaces in order to show how our new definition of Bayesianism implies the structure of Bayesian belief spaces.

3. Bayesian belief spaces

There are two requirements on a belief space that make it a model of a Bayesian agent. Namely, that the agent is certain of his type and that his posterior belief is generated by conditioning his prior belief on his type. In the next two subsections we discuss and formulate these two requirements, and in the last subsection we formally define Bayesian belief spaces.

3.1. BEING CERTAIN OF ONE'S TYPE

We now express formally the phrase "the agent is *certain of (*or *knows) his type*." As the agent is certain of *events*, we need, first, to describe "his type" by an event. We denote by $T(\omega)$ the set of all the states in which the agent's type is $t(\omega)$, that is,

$$T(\omega) = \{\omega' \mid t(\omega') = t(\omega)\}. \tag{1}$$

The set $T(\omega)$ is the natural candidate for the event "the agent's type is $t(\omega)$," but we still need to show that this set is indeed an event, that is, a measurable set. This follows from the next proposition.

PROPOSITION 1: *for each state ω,*

$$T(\omega) = \bigcap_{\{(p,E) \mid \omega \in B^p(E)\}} B^p(E), \tag{2}$$

where the intersection is taken over all pairs of real number p and $E \in \Sigma$, that satisfy the required condition. This equality also holds when the intersection is restricted to pairs of rational p and $E \in \Sigma_0$, *that satisfy the condition. Therefore, as a countable intersection of events, $T(\omega)$ is an event.*

In addition to showing that $T(\omega)$ is an event, this proposition suggests an alternative description of the event that the agent is of type $t(\omega)$: it is the intersection of all the events that describe the agent's belief at ω.

Now, with $T(\omega)$ being the event that the agent's type is his type at ω, we can define what it means for the agent to be certain of his type in a given state.

DEFINITION 2: *the agent is said to be* certain of his type *at ω when*

$$t(\omega)(T(\omega)) = 1. \tag{3}$$

Let us denote by C the set of all states in which the agent is certain of his type, that is,

$$C = \{\omega \mid t(\omega)(T(\omega)) = 1\}.$$

We show that C is a measurable set, and thus it is the event that the agent is certain of his type. This follows from the next proposition, which is used also in the proof of the main theorem. We denote by $\neg X$ the complement of the event X.

PROPOSITION 2: *the set C of all states in which the agent is certain of his type satisfies*

$$C = \bigcap_{p,E} \left(\neg B^p(E) \cup B^1(B^p(E))\right) \tag{4}$$

where the intersection is taken, either over all p and $E \in \Sigma$, or over rational p and $E \in \Sigma_0$. Thus, C is an event.

We can now define the condition that the agent is certain of his type, which is the first requirement that Bayesian spaces should satisfy.

DEFINITION 3: *the agent is* certain of his type *in a belief space with a prior μ, if he is certain of his type almost everywhere with respect to μ, that is, if $\mu(C) = 1$.*

Note that $\neg B^p(E) \cup B^1(B^p(E))$ is the event that either the probability ascribed by the agent to E is less than p, or else, if it is at least p, then the agent is certain that the probability he ascribes to E is at least p. In short, this is the event that *if* the probability he ascribes to E is at least p *then* he is certain of this fact. Proposition 2 says that the event that the agent is certain of his type is the event that he is certain of any belief of his that he holds.

3.2. CONDITIONING ON ONE'S TYPE

Before we express formally the requirement that the posterior of the agent is formed by conditioning the prior on his type, we discuss the simple case in which the state space Ω is countable. In this case the requirement of conditioning says that in each state ω, for which the prior probability of the event $T(\omega)$ is positive, and for all events E,

$$t(\omega)(E) = \mu\left(E \mid T(\omega)\right). \tag{5}$$

But (5) cannot serve as a general definition for all belief spaces, including the uncountable ones, as in such spaces $\mu(T(\omega))$ may be zero for a set of ω's of positive μ-probability. Therefore we adopt a non-local definition (i.e., one which is not defined, like (5), for a single state) of conditioning on type. We first describe the logic of this definition for the countable case.

Multiplying both sides of (5) by $\mu(\omega)$, and summing over states for which the right hand side is defined, we have,

$$\sum_{\omega} t(\omega)(E)\mu(\omega) = \sum_{\omega} \mu\left(E \mid T(\omega)\right)\mu(\omega). \tag{6}$$

After grouping terms in the right hand side of (6), corresponding to states in the same $T(\omega)$, it can be rewritten as $\sum \mu(E \mid T(\omega))\mu(T(\omega))$, where the summation is over the $T(\omega)$'s. By the Theorem of Total Probability, this is just $\mu(E)$. Thus, if (5) is satisfied for μ-almost all ω, then,

$$\mu(E) = \sum_{\omega} t(\omega)(E)\mu(\omega). \tag{7}$$

It is easy to see that when the agent is certain of his type, then (7) implies that (5) holds for μ-almost all ω. Indeed, suppose that (3) holds for ω for which $\mu(\omega) > 0$. By (7), for any E, $\mu(E \cap T(\omega)) = \Sigma_{\omega'} t(\omega')(E \cap T(\omega))\mu(\omega')$. By (3), this sum is $\Sigma_{\omega' \in T(\omega)} t(\omega')(E)\mu(\omega')$. By (1), all the terms are constantly $t(\omega)(E)$, and therefore the latter sum is $t(\omega)(E)\mu(T(\omega))$, which is (5).

Condition (7), unlike (5), is a global condition which is easily generalized to all belief spaces in the following definition.

DEFINITION 4: *we say that the prior μ is* invariant *if for each event E,*

$$\mu(E) = \int t(\omega)(E)\, d\mu(\omega). \tag{8}$$

Observe that the invariance of μ does not imply that the agent is certain of his type, as the following example demonstrates.

EXAMPLE 1: let $\Omega = \{\omega_1, \omega_2\}$, $t(\omega_1) = (1/3, 2/3)$, $t(\omega_2) = (2/3, 1/3)$, and $\mu = (1/2, 1/2)$. Then, μ is invariant but the agent is not certain of his type in either state, as $t(\omega_i)(T(\omega_i)) = 1/3$, for $i = 1, 2$.

Thus, in order that invariance implies conditioning on the agent type, we have to assume also that the agent is certain of his type, as explained above for the countable case.

The term *invariance* was chosen to describe a prior that satisfies (8), because formally the function t can be thought of as a transition function of a Markov chain on Ω. Formula (8) says that μ is an invariant probability measure of this Markov chain. (See Samet, 1998*a*, for implications of this observation.) Mertens and Zamir (1985) and Feinberg (1996) also use (8) to describe Bayesian priors.

3.3. DEFINING BAYESIAN BELIEF SPACES

We are ready now to define Bayesian spaces in terms of the two properties discussed above.

DEFINITION 5: *a belief space is* Bayesian *if*

(1) The agent is certain of his type.
(2) The prior is invariant.

Bayesian agents are described, in the economic and game theoretic literature, almost exclusively by using Bayesian spaces, as defined here—except, perhaps, for small variations and differences in formulation.

4. Characterizing Bayesian agents

Our main result provides a characterization of Bayesian spaces which does not make use of properties 1 and 2 in definition 5. We show that a belief space is Bayesian if and only if whenever the prior is conditioned on some specification of the posterior belief, it agrees with this specification. A posterior belief is specified, in this theorem, by events of the form $\cap_{k=1}^{n} B^{p_k}(E_k)$, which put bounds on the probability of finitely many events. Events like this can approximate the posterior probability of any finite list of events to any accuracy, as $B^p(E) \cap B^q(\neg E)$ bounds the probability of E to be in the interval $[p, 1-q]$.

THEOREM 1: *a belief space with a prior μ is Bayesian iff for each $n \geq 1$, all events $E_1, \ldots, E_n$, and all numbers $p_1, \ldots, p_n$,*

$$\mu\left(E_m \mid \bigcap_{k=1}^{n} B^{p_k}(E_k)\right) \geq p_m, \tag{9}$$

for $1 \leq m \leq n$, whenever this conditional probability is defined.

No assumption is made in this condition that the agent acquires any information with certainty. That is, (9) does not depend on there being some non-trivial event of which the posterior probability is 1. It suggests an interpretation of Bayesianism as a simple requirement of consistency between prior and posterior, rather than a description of the mechanism by which the posterior belief is derived from a prior one, namely conditioning on acquired certain information. Yet, (9) implies, by theorem 1, that the agent is certain of his type. Being so, according to the suggested interpretation, does not represent the learning with certainty, of any new facts (*signals*, in the vernacular of economists) which initiate the updating of the prior. Rather, certainty of type reflects the capacity of the agent to be certain of his own (posterior) belief, (in concert with the presentation of $T(\omega)$ in proposition 2) and as such it is the *result* of the updating, and not its cause.

The following theorem states that even with the restriction $n = 2$, the condition (9) is sufficient for a belief space to be Bayesian.

THEOREM 2: *a belief space with a prior μ is Bayesian iff for any two events E and F, and numbers p and q,*

$$\mu(E \mid B^p(E) \cap B^q(F)) \geq p, \tag{10}$$

whenever this conditional probability is defined.

Next, we consider a restriction of (10) to conditioning events of the form $B^p(E) \cap B^q(\neg E)$. Conditioning both E and $\neg E$ on this event, yields by (10),

$$1 - q \geq \mu(E \mid B^p(E) \cap B^q(\neg E)) \geq p. \tag{11}$$

This is a natural requirement, as the probability of E is conditioned in (11) on an event which concerns just the posterior belief on E. The stronger requirements, (9) and (10), allow including in the conditioning event information about the posterior of events other than E. These conditions add to (11) the requirement that the only part of the posterior belief which is relevant as condition for the prior of E is the posterior probability of E.

It turns out that the weaker condition (11) is not sufficient to guarantee that a belief space is Bayesian. We can still state the following implication of (11).

PROPOSITION 3: *if in a belief space with a prior μ, (11) holds for any event E and numbers p and q, whenever this conditional probability is defined, then μ is invariant.*

The following example demonstrates that (11) does not imply that a belief space is Bayesian.

EXAMPLE 2: let $\Omega = \{\omega_1, \omega_2, \omega_3\}$. The agent's type function is: $t(\omega_1) = (0, 1/2, 1/2), t(\omega_2) = (1/2, 0, 1/2)$ and $t(\omega_3) = (1/2, 1/2, 0)$, and the prior is $\mu = (1/3, 1/3, 1/3)$. It is easy to check that (11) holds for all E, p and q. Yet, this space is clearly not Bayesian, as the agent is not certain of his type in any of the states.

5. Proofs

PROOF OF PROPOSITION 1

Suppose $\omega' \in T(\omega)$. Then, for any E and p, $\omega \in B^p(E)$ iff $\omega' \in B^p(E)$ and therefore ω' is in the set on the right hand side of (2). Conversely, suppose ω' is in the set on the right hand side of (2). Fix $E \in \Sigma_0$. Then, whenever ω is in $B^p(E) \cap B^{1-q}(\neg E)$ for some rational $p < q$, ω' is also in this event. Thus $t(\omega)(E)$ and $t(\omega')(E)$ belong to the same intervals $[p, q]$ with rational ends. But this implies that $t(\omega)(E) = t(\omega')(E)$. Hence, $t(\omega)$ and $t(\omega')$ are two probability measures that agree on Σ_0 and therefore the same probability measure. Hence $\omega' \in T(\omega)$. ■

PROOF OF PROPOSITION 2

Suppose the agent is certain of his type in state ω, that is, $\omega \in C$. Fix E and p. Then, either $\omega \in \neg B^p(E)$, or else, by proposition 1, $t(\omega)(B^p(E)) = 1$, that is $\omega \in B^1(B^p(E))$. Thus, $\omega \in \neg B^p(E) \cup B^1(B^p(E))$.

Conversely, suppose that ω belongs to the event in the right hand side of (4), where the intersection is taken over rational p and events $E \in \Sigma_0$. Then, for each such p and E, if $\omega \in B^p(E)$, then $\omega \in B^1(B^p(E))$. This means that

$$\omega \in \bigcap_{\{(p,E)|\omega \in B^p(E)\}} B^1\left(B^p(E)\right),$$

where the intersection is over rational p and $E \in \Sigma_0$ that satisfy the required condition. But B^1 commutes with countable intersections and thus,

$$\omega \in B^1\left(\bigcap_{\{(p,E)|\omega\in B^p(E)\}} B^p(E)\right).$$

By proposition 1, this is equivalent to saying that $\omega \in B^1(T(\omega))$. Therefore $t(\omega)(T(\omega)) = 1$, and thus $\omega \in C$. ■

PROOF OF PROPOSITION 3

To simplify the notation, we introduce an abbreviation for the event that the agent's posterior probability of E is in the interval $[p,q)$, as follows:

$$B^{p,q}(E) = B^p(E) \setminus B^q(E).$$

We assume that (11) holds, and show that for any given $\varepsilon > 0$ the two sides of (8) can differ by ε at most.

Choose a sequence of numbers $0 = p_0 < p_1 < \cdots < p_m > 1$, such that $|p_k - p_{k-1}| \le \varepsilon$ for $k = 1,\ldots,m$. Clearly, the events $B^{p_{k-1},p_k}(E)$ for $k = 1,\ldots,m$ form a partition of Ω and therefore,

$$\mu(E) = \sum_{k=1}^{m} \mu\left(E \cap B^{p_{k-1},p_k}(E)\right), \tag{12}$$

and

$$\int t(\omega)(E)\,d\mu(\omega) = \sum_{k=1}^{m} \int_{B^{p_{k-1},p_k}(E)} t(\omega)(E)\,d\mu(\omega). \tag{13}$$

To evaluate the k term in (12) we consider an event $F^r = B^{p_{k-1}}(E) \cap B^{1-r}(\neg E)$ for $p_{k-1} < r < p_k$. This is the event that the posterior probability of E is in the interval $[p_{k-1}, r]$. By (11), $r\mu(F^r) \ge \mu(E \cap F^r) \ge p_{k-1}\mu(F^r)$. When r converges to p_k the events F^r converge monotonically to $B^{p_{k-1},p_k}(E)$, and hence, by the continuity of μ, we have,

$$p_{k-1}\mu\left(B^{p_{k-1},p_k}(E)\right) \le \mu\left(E \cap B^{p_{k-1},p_k}(E)\right) \le p_k\mu(B^{p_{k-1},p_k}(E))$$

To evaluate the k term in (13) we note that for $\omega \in B^{p_{k-1},p_k}(E)$, $t(\omega)(E) \in [p_{k-1}, p_k)$, and thus

$$p_{k-1}\mu\left(B^{p_{k-1},p_k}(E)\right) \le \int_{B^{p_{k-1},p_k}(E)} t(\omega)(E)\,d\mu(\omega) \le p_k\mu\left(B_i^{p_{k-1},p_k}(E)\right).$$

Therefore, the k terms in (12) and (13) have the same bounds and they can differ by at most the difference between these bounds, which is bounded by $\varepsilon\mu(B^{p_{k-1},p_k}(E))$. Thus, by (12) and (13), the two sides of (8) can differ by ε at most. ■

PROOF OF THEOREMS 1 AND 2

Suppose that the prior μ satisfies (10). By proposition 3, μ is an invariant prior, and therefore to show that the space is Bayesian, it is enough to prove that the agent is certain of his type. Assume, to the contrary, that he is not. Then the μ-probability of the complement of C is positive. By proposition 2, this means that,

$$\mu\left(\bigcup_{p,E} B^p(E) \cap \neg B^1(B^p(E))\right) > 0,$$

where the union is countable. Therefore, for some p and E,

$$\mu\left(B^p(E) \cap \neg B^1(B^p(E))\right) > 0.$$

But $\neg B^1(B^p(E))$—the event that the agent is not certain of $B^p(E)$—is the same as the event that he ascribes positive probability to $\neg B^p(E)$. Hence $\neg B^1(B^p(E)) = \cup_{r>0} B^r(\neg B^p(E))$, where r is rational. Thus, for some $r > 0$,

$$\mu\left(B^p(E)\right) \cap B^r\left(\neg B^p(E)\right) > 0 \tag{14}$$

Applying (10), as the conditional probability is well defined by (14), yields

$$\mu\left(\neg B^p(E)\right) \mid B^p(E) \cap B^r\left(\neg B^p(E)\right) \geq r > 0.$$

But this is a contradiction, since $\neg B^p(E)$ and the conditioning event are disjoint.

Conversely, suppose that μ is a prior in a Bayesian belief space. We show that (9) holds. Applying the invariance condition (8) to $\cap_{k=1}^n B^{p_k}(E_k) \cap E_m$, we have

$$\mu\left(\bigcap_{k=1}^n B^{p_k}(E_k) \cap E_m\right) = \int t(\omega)\left(\bigcap_{k=1}^n B^{p_k}(E_k) \cap E_m\right) d\mu(\omega). \tag{15}$$

Now, if $\omega \in \cap_{k=1}^n B_i^{p_k}(E_k)$, then $T(\omega) \subseteq \cap_{k=1}^n B_i^{p_k}(E_k)$, and otherwise, $T(\omega) \subseteq \neg \cap_{k=1}^n B_i^{p_k}(E_k)$. Hence, if the agent knows his type at ω, then $t(\omega)(\cap_{k=1}^n B^{p_k}(E_k))$ is 1 in the first case, and 0 in the other. Therefore, $t(\omega)(\cap_{k=1}^n B^{p_k}(E_k) \cap E_m) = t_i(\omega)(E_m) \geq p_m$, in the first case, and $t(\omega)(\cap_{k=1}^n B^{p_k}(E_k) \cap E_m) = 0$, in the other.

As the agent knows his type μ-almost everywhere, we can rewrite (15) as,

$$\mu\left(\bigcap_{k=1}^n B^{p_k}(E_k) \cap E_m\right) = \int_{\cap_{k=1}^n B^{p_k}(E_k) \cap E_m} t(\omega)(E_m)\, d\mu(\omega).$$

Since the integrand is no less than p_m, we conclude that,

$$\mu\left(\bigcap_{k=1}^{n} B^{p_k}(E_k) \cap E_m\right) \geq p_m \mu\left(\bigcap_{k=1}^{m} B^{p_k}(E_k)\right),$$

which proves (9). ■

References

Battigalli, P. & Bonanno, G. (1997). The logic of belief persistence. *Economics and Philosophy*, **13**, 39–59.

Feinberg, Y. (1996). *Characterizing the existence of a common prior via the notion of disagreement*. Doctoral dissertation, The Hebrew University, Jerusalem.

Gaifman, H. (1988). A theory of higher order probabilities. In B. Skyrms & W. Harper, Eds. *Causation, Chance and Credence*. Kluwer.

Halpern, J.Y. (1991). The relationship between knowledge, belief, and certainty. *Annals of Artificial Intelligence and Mathematics*, **4**, 301–322.

Halpern, J.Y. (1998). Errata: The relationship between knowledge, belief, and certainty. *Annals of Artificial Intelligence and Mathematics*, page forthcoming.

Harsanyi, J.C. (1967–68). Games with incomplete information played by bayesian players, part i, ii, & iii. *Management Science*, **14**, 159–182, 320–334, 486–502.

Heifetz, A. & Samet, D. (1998). Topology-free typology of beliefs. *Journals of Economic Theory*, **82**, 263–293.

Lewis, D. (1980). A subjective guide to objective chance. In R. C. Jeffrey, Ed. *Studies in Inductive Logic and Probability II*, pp. 263–293. California: University of California Press.

Maher, P. (1993). *Betting on Theories.* Cambridge: Cambridge University Press.

Mertens, J.F. & Zamir, S. (1985). Formulation of bayesian analysis for games with incomplete information. *International Journals of Game Theory*, **14**, 1–29.

Miller, D. (1966). A paradox of information. *British Journal for the Philosophy of Science*, **17**, 59–61.

Popper, K.R. (1966). A comment on Miller's new paradox of information. *British Journal for the Philosophy of Science*, **17**, 61–69.

Samet, D. (1997). On the triviality of high-order beliefs. *Journal of Philosophical Logic*, page forthcoming.

Samet, D. (1998*a*). Common priors and iterated expectations. *Games and Economic Behavior*, **24**, 131–141.

Samet, D. (1998*b*). Quantified beliefs and believed quantities. In I. Gilboa, Ed. *Theoretical Aspects of Rationality and Knowledge: Proc. Seventh Conference*. Morgan Kaufmamn.

Skyrms, B. (1980). *Causal Necessity*. Yale: Yale University Press.

van Fraassen, B. (1984). Belief and the will. *Journal of Philosophy*, **81**, 235–256.