LONG-RUN IMPLICATIONS OF MAXIMIZING POSTERIOR EXPECTED UTILITY

FATEMEH BORHANI AND EDWARD J. GREEN
PENNSYLVANIA STATE UNIVERSITY

1. Introduction

Let Ω be the set of states. For any $t \in \mathbb{N}$ consider finite partition Π_t on Ω such that Π_{t+1} be weakly finer than Π_t. Let $\mathcal{F}_t$ be the algebra of events generated by Π_t and let $\Delta(\Omega)$ be the set of probability measures on Ω. $P \in \Delta(\Omega)$ will be called a belief. One aspect of Bayesian theory is that information will never be thrown away. A Bayesian agent updates her beliefs using all her information. Can a contingent plan which only depends on the most recent piece of information be rationalized even if it is not a sufficient statistic for the entire history? As it will be shown in this paper, the answer is yes.

Another property of Bayesian theory is that posterior beliefs converge almost surely to truth.[1] How about Bayesian behavior? How does action choice of a Bayesian agent changes as t increases and information becomes more precise? Does her choice asymptotically converge as her beliefs converge? The example in next section will show that the answer is: Not necessarily so.

Then the main question of this paper is: What asymptotic properties should a Bayesian contingent plan have? We argue that neither asymptotically converging nor depending on all the information are necessary. The goal of present work is to show there are no such intrinsic asymptotic properties. Any contingent plan which satisfies a consistency property, which is an obvious revealed-preference implication of the sure-thing principle together with assuming positive probability of observing any finite initial sequence of observations, and will be defined later, can be rationalized.

The decision problem that we have in mind is a one-shot decision making. When we talk about a particular point of time, the only relevant factor is the finite amount of information that has been observed up to that point. At any point of time, contingent plan specifies the action that agent would choose given the finite amount of information observed up to that point.

2. Example: a chaotic contingent plan

Let state space be $\Omega = \{1, 2\}^{\mathbb{N}}$ and action set be $A = \{0, 1, 2\}$. We show a typical element in Ω by $\omega = (\omega_1, \omega_2, ...)$. Define $\omega^t = (\omega_1, \omega_2, ..., \omega_t)$ to be finite sequence of observations from date 1 to date t. Each state has a unique representation in the graph shown in figure 1.

Date: 2014.07.05.

The authors gratefully acknowledge the Human Capital Foundation, (`http://www.hcfoundation.ru/en/`) and particularly Andrey P. Vavilov, for research support through the Center for the Study of Auctions, Procurements, and Competition Policy (CAPCP, `http://capcp.psu.edu/`) at the Pennsylvania State University. Green appreciates the hospitality of CEMFI, where part of his research was conducted.

[1]For example see Breiman [1992] theorem 5.21

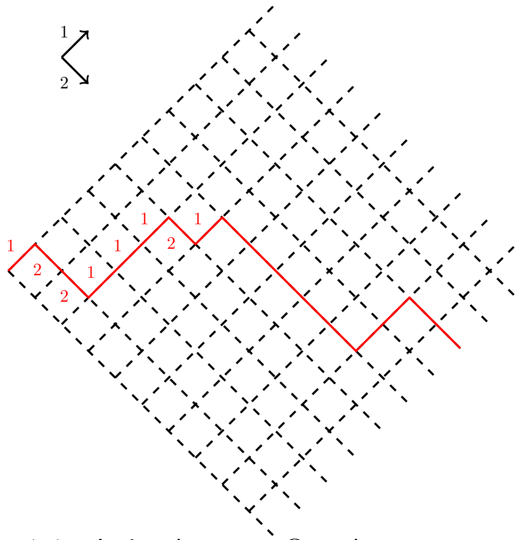

FIGURE 1. There is a 1-1 relation between Ω and representateions on this graph. For example $\omega^8 = (1, 2, 2, 1, 1, 1, 2, 1)$

Consider contingent plan $\beta(\omega^t) = \sum_{t' \leq t} \omega_{t'} \pmod 3$. Let $\beta(\emptyset) = 0$. Figure 2 represents actions suggested by this contingent plan. This contingent plan has the property that, at every ω and t, $\beta(\omega^t)$, $\beta(\omega^t * \{1\})$ and $\beta(\omega^t * \{2\})$ are distinct values from one another. Therefore, for any ω, chosen actions never settle down.

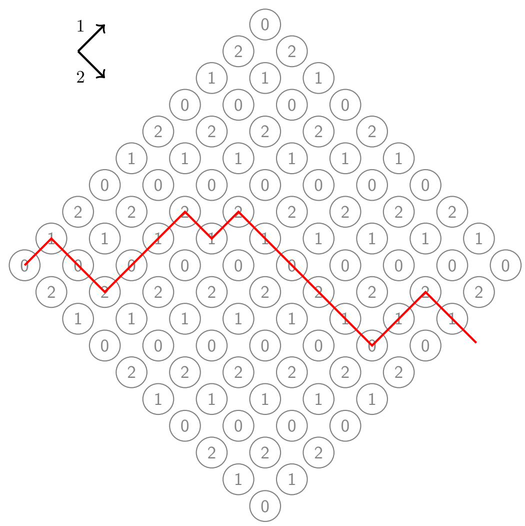

FIGURE 2. Suggested actions by β

Intuitively, such a contingent plan seems not to reflect Bayesian decision making. Nevertheless, it will be shown here that it can be rationalized for some state dependent utility function and some probability measure over state space. But before we prove this claim, we need some definitions.

Let $B^*(\omega^t)$ be the cylinder set defined as

$$B^*(\omega^t) = \{\zeta \in \Omega | \zeta^t = \omega^t\}. \tag{1}$$

Define Ψ to be the set of states that after some finite time zigzag forever

$$\Psi = \{\omega \in \Omega | \exists t \in \mathbb{N} \quad \forall t' \geq t \quad \omega_{t'} \neq \omega_{t'+1}\} \tag{2}$$

and define $\tau \colon \Omega \to \mathbb{N}$ be such that for $\omega \in \Psi$ it indicates the initiation of zigzagging

$$\tau(\omega) = \min\{t | \forall t' > t \quad \omega_{t'} \neq \omega_{t'+1}\}. \tag{3}$$

Observe that given some $t > 0$ and given some ω there exists only one state in Ψ consistent with ω^t for which τ is equal to t. Let set $C(\psi)$ be such that for $\psi \in \Psi$, it consists of the two actions that are picked by β at time $\tau(\psi)$ and later:

$$C(\psi) = \{\beta(\omega^{\tau(\omega)}), \beta(\omega^{\tau(\omega)+1})\} = \{\sum_{t \leq \tau(\omega)} \omega_t \pmod 3, \sum_{t \leq \tau(\omega)+1} \omega_t \pmod 3\}. \tag{4}$$

Define $\hat{\Psi}(\omega^t)$ to be the set of states in Ψ with the same initial t signals as ω that begin to zigzag at t or before it:

$$\hat{\Psi}(\omega^t) = \{\psi \in \Psi \cap \mathcal{B}^*(\omega^t) | \tau(\psi) \leq t\}. \tag{5}$$

Now we show that for any t and any ω,

$$\bigcap_{\eta \in \hat{\Psi}(\omega^t)} C(\zeta) = \{\beta(\omega^t)\}. \tag{6}$$

First consider $\emptyset$. $\hat{\Psi}(\emptyset) = \{(1,2,1,2,\dots),(2,1,2,1,\dots)\}$. Considering that $C((1,2,1,2,\dots)) = \{0,1\}$ and $C((2,1,2,1,\dots)) = \{2,0\}$ we can see that (6) is satisfied at $\emptyset$. Now consider some $t > 0$ and some ω. $\hat{\Psi}(\omega^t)$ has two elements one which starts zigzagging at t, lets call it ζ and the other at $t-1$ or earlier, lets call it ξ. Then $C(\zeta) = \{\beta(\omega^t), \beta(\omega^{t+1})\}$ and $C(\xi) = \{\beta(\omega^t), \beta(\omega^{t-1})\}$. Therefore (6) is satisfied.

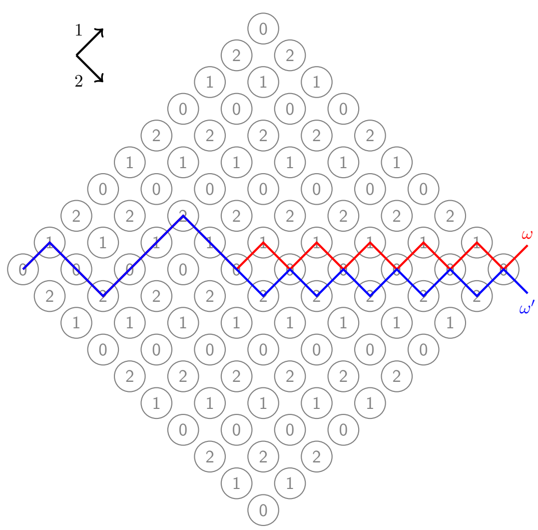

FIGURE 3. $\tau(\omega) = 7, C(\omega) = \{0,1\}, \tau(\omega') = 8, C(\omega') = \{0,2\}$

Now pick $\delta \in (0, \frac{1}{2})$ and let $\kappa = \frac{1}{1-2\delta}$ Consider function

$$P(\{\omega\}) = \kappa \delta^{\tau(\omega)} 1_{\Psi}(\omega). \tag{7}$$

Because of $\sum_{\psi \in \Psi} P(\psi) = \sum_{t \geq 0} \kappa \delta^t \cdot 2^t = 1$, this function defines a countably additive probability measure on Ω.[2] Since $\delta < \frac{1}{2}$, $\delta^t > \frac{\delta^{t+1}}{1-\delta}$ so

$$\min\{P(\{\psi\})|\psi \in \hat{\Psi}_t(\omega)\} = \kappa\delta^t > \kappa\frac{\delta^{t+1}}{1-\delta} = P(\Psi \cap B^*(\omega^t) \setminus \hat{\Psi}_t(\omega)). \tag{8}$$

Consider state dependent utility function which for $\psi \in \Psi$ satisfies

$$u(a, \psi) = 1_{C(\psi)}(a). \tag{9}$$

When the probability measure and utility function are defined in this way, expected net gain from any deviation from β is negative. To see this, note that regarding (6), $\beta(\omega^t)$ is the only action which attains maximum state dependent utility for all the states in $\hat{\Psi}_t(\omega)$. Therefore the loss of deviation is at least as large as $\min\{P(\psi) \cdot u(\beta(\omega^t), \psi)|\psi \in \hat{\Psi}_t(\omega)\} = \min\{P(\psi)|\psi \in \hat{\Psi}_t(\omega)\}$. Maximum expected gain is when $\beta(\omega^t)$ attains minimum state dependent utility for all states outside $\hat{\Psi}_t(\omega)$ while deviation action attains maximum state dependent utility for all these states.Therefore the expected gain is at most equal to $\sum_{\psi \in \Psi \cap B^*(\omega^t) \setminus \hat{\Psi}_t(\omega)} P(\psi)\{u(a, \psi) - u(\beta(\omega^t), \psi)\} = \sum_{\psi \in \Psi \cap B^*(\omega^t) \setminus \hat{\Psi}_t(\omega)} P(\psi) = P(\Psi \cap B^*(\omega^t) \setminus \hat{\Psi}_t(\omega))$. But (8) shows that expected loss is strictly greater that expected gain and therefore expected net gain is negative. Thus β is the unique maximizing contingent plan for our choice of state dependent utility function and probability measure.

3. Observation sequences and contingent plan

Let A be finite set of actions and X be finite set of signals. Define the set of *finite observation sequences*, $\boldsymbol{X}^* = \bigcup_{t \in \mathbb{N}} \boldsymbol{X}^t$. We show a typical element of $\boldsymbol{X}^*$ by $\vec{x}$. Let $\emptyset$ denote the null sequence and let $\vec{x} \preceq \vec{y}$ denote that $\vec{y}$ is an extension of $\vec{x}$ and let $\vec{x} \prec \vec{y}$ denote that $\vec{y}$ is a proper extension of $\vec{x}$. If $\vec{x} = (x_1, \ldots, x_n)$, then for $1 \leq t \leq n$, $\vec{x}_t$ denotes x_t. An *immediate extension* of $\vec{x} \in \boldsymbol{X}^*$ is $\vec{x}$ followed by a single observation. Let $\vec{x} * y$ denote the immediate extension of $\vec{x}$ by y. Function $\alpha\colon \boldsymbol{X}^* \to A$ is called a contingent plan.

The usual representation of observations in terms of measurable functions is adopted here. A measurable space $(\Omega, \mathcal{B})$ is constructed from $\boldsymbol{X}$ by taking $\Omega = \boldsymbol{X}^{\mathbb{N}_+}$ and projection functions $X_t(\omega) = \omega_t$, and by taking $\mathcal{B}$ to be the smallest σ-algebra with which all of the projection functions X_t are measurable. $\mathcal{B}_t \subseteq \mathcal{B}$ is the σ-algebra generated by $X_1, \ldots, X_t$. In particular, every cylender set $B^*(\vec{x}) = \{\omega | \forall t \leq \lambda(\vec{x}) \;\; X_t(\omega) = \vec{x}_t\}$ is measurable with respect to $\mathcal{B}$ and to $\mathcal{B}_t$ for $\vec{x} \in \boldsymbol{X}^{t'} \quad t' \leq t$. Let $\vec{x} \preceq \omega$ denote that $\omega^t = \vec{x}$.

4. Consistent contingent plans

Define contingent plan α to be *consistent* at $\vec{x}$ iff

$$\forall a \in A \setminus \{\, \alpha(\vec{x})\} \quad \exists y \in X \quad \alpha(\vec{x} * y) \neq a. \tag{10}$$

Consistency is an obvious revealed-preference implication of the sure-thing principle formulated by Savage [1954][3] together with the assumption of positive probability for any initial finite sequence of observations. Therefore, a Bayesian agent must be consistent. If there is a finite bound on the length of observation sequences, then the converse is true: a finite contingent plan that is consistent at every node must be Bayes for some u and P. Green and Park [1996] prove this fact, in the course of proving theorem 1 in that paper. (Antecedent investigations of the relationship between being consistent and

[2] The factor 2^t in the sum of probabilities reflects there being 2^t sequences of length t, each of which corresponds to an element of Ψ that is assigned probability $\kappa\delta^t$.

[3] In the formal theory introduced by Savage, which has been adopted by most subsequent researchers, acts are defined entities. Despite that difference from the present theory (in which A is a primitive of the theory), clear parallels can be drawn.

being Bayes include Weller [1978], Hammond [1988], and Epstein and Breton [1993].) They also state an infinite-horizon result However, the theorem, which concerns infinite-horizon plans, is false. (Its proof contains an invalid assertion that an a.s. convergent martingale also converges in L_1.)

Proposition 6 in the present paper, to be stated and proved below, is a sound weakening of Green and Park's theorem 1. The proof of this proposition generalize the preceding example. In the first step we define Ψ a subset of Ω which has three properties. Firstly, it is countable. Secondly, for any finite sequence of observations, there exists some $\psi \in \Psi$ which is not in contradiction with observations. And lastly, for any $\psi \in \Psi$, we can define an index $\tau(\psi)$ and a set of actions $C(\psi, \tau(\psi))$ such that $\alpha(\psi^t)$ be the unique element in common in $C(\psi, \tau(\psi))$ for any $\psi \in \Psi$ consistent with observations and with an index smaller or equal to t. Consistency condition is used in lemma 1. Lemma 1 to lemma 4 are used in defining Ψ and proving these properties.

In the next step of proof, we show that mentioned properties of Ψ, enable us to choose a bounded state dependent utility function (defined in terms of C) and a probability measure (defined in terms of τ) with full support on Ψ such that β is the unique expected utility maximizing contingent plan. Note that, since every finite sequence of observations is an initial sequence at some state in Ψ, every finite initial sequence has positive probability of being observed.

5. Step 1: Defining $\Psi \subseteq \Omega$

Define

$$C(\omega, t) = \{\ \alpha(\omega^{t+n}) | n \in \mathbb{N}_+\} \tag{11}$$

and

$$D(a, \omega^t) = \{\psi \in \Omega | \omega^t \prec \psi, a \neq C(\psi, t)\}\ . \tag{12}$$

Lemma 1. *If α is consistent and $a \neq \alpha(\omega^t)$, then $D(a, \omega^t) \neq \emptyset$.*

Proof. A recursive construction, using consistency at each stage $n > 0$, establishes that there is a sequence $\vec{x}^0, \vec{x}^1, \ldots$ such that $\vec{x}^0 = \omega^t$ and, for every $n \in \mathbb{N}$, $a \neq \alpha(\vec{x}^n)$ and $\vec{x}^{n+1}$ is an immediate proper extension of $\vec{x}^n$. Then, defining $\omega_t = x \iff \exists n\ \vec{x}^n_t = x$ produces an element of Ω that satisfies the lemma. □

The next lemma follows immediately from lemma 1

Lemma 2. *If α is consistent, then there exists a function $d\colon A \times \boldsymbol{X}^* \to \Omega \cup \{\emptyset\}$ such that, for every $\vec{x}$,*

$$d(\alpha(\vec{x}), \vec{x}) = \emptyset \text{ and } \forall a \neq \alpha(\vec{x})\ \ d(a, \vec{x}) \in D(a, \vec{x})\,. \tag{13}$$

Define $\Psi \subseteq \Omega$ by

$$\Psi = \{d(\alpha(\vec{x}), \vec{x}) \mid a \neq \alpha(\vec{x})\} \tag{14}$$

Define $\tau : \Omega \to \mathbb{N}$ such that for any $\psi \in \Psi$ it satisfies

$$\tau(\psi) = \min\{\ t \in \mathbb{N} | \text{for some } a \neq \alpha(\psi^t),\ a \notin C(\psi^t)\}. \tag{15}$$

Define $\hat{\Psi}(\omega^t) \subseteq B^*(\omega^t) \cap \Psi$ by

$$\hat{\Psi}(\omega^t) = \{\psi \in \Psi \cap B^*(\omega^t) | \tau(\psi) \leq t\}. \tag{16}$$

Note that equivalently we can restate $\hat{\Psi}(\omega^t)$ as

$$\hat{\Psi}(\omega^t) = \bigcup_{t' \leq t} \bigcup_{a \neq \alpha(\omega^{t'})} \{\ d(a, \omega^{t'})\}. \tag{17}$$

therefore $\hat{\Psi}(\omega^t)$ is finite.

Lemma 3. Ψ *is countable.*

Proof. Ψ is a subset of the range of a function with a countable domain therefore it is countable. □

Lemma 4.

$$\bigcap_{\psi \in \hat{\Psi}(\omega^t)} C(\psi, \tau(\psi)) = \alpha(\omega^t) \tag{18}$$

Proof. Please note that

$$\forall \psi \in \hat{\Psi}(\omega^t) \quad \alpha(\omega^t) \in C(\psi, \tau(\psi)) \tag{19}$$

and

$$\bigcap_{\{\psi \in \hat{\Psi}(\omega^t) | \tau(\psi) = t\}} C(\psi, \tau(\psi)) \;=\; \alpha(\omega^t). \tag{20}$$

Result immediately follows from the above two equations. □

6. Step 2: Probability Measure and Utility Function

Since Ψ is countable and every $\hat{\Psi}(\omega^t)$ is nonempty and finite, there exists a countably additive probability measure P such that

$$P(\Psi) = 1, \tag{21}$$

$$\forall \psi \in \Psi \; P(\{\psi\}) > 0 \tag{22}$$

and

$$\sum_{\zeta \in \Psi \cap B^*(\omega^t) \setminus \hat{\Psi}(\omega^t)} P(\{\zeta\}) = \delta \times \min_{\hat{\Psi}(\omega^t)} \{P(\{\omega\})\} > 0 \tag{23}$$

for some $\delta \in (0, 1)$. Note that conditions (22) entails that every finite sequence of initial observations has positive probability. Define bounded utility function

$$u(a, \omega) = \begin{cases} 1 & \text{if } a \in C(\psi, \tau(\psi)) \text{ and } \omega \in \Psi \\ 0 & \text{otherwise.} \end{cases} \tag{24}$$

Lemma 5. *If P and u satisfy conditions (22) and (24), and $a \neq \alpha(\omega, t)$, then*

$$U(a, \omega, t) < U(\alpha(\omega^t), \omega, t)\,. \tag{25}$$

where U is derived from u according to

$$U(a, \omega, t) = \frac{\int_{B^*(\omega^t)} u(a, \psi)\, dP(\psi)}{B^*(\omega^t)} \tag{26}$$

Proof. An upper bound for the maximum gain of deviating from $\alpha(\omega^t)$ is

$$\sum_{\zeta \in \Psi \cap B^*(\omega^t) \setminus \hat{\Psi}(\omega^t)} P(\{\zeta\}) \times 1 = \delta \times \min_{\zeta \in \hat{\Psi}(\omega^t)} \{P(\{\zeta\})\} \tag{27}$$

Where the equality follows from (23). With regard to lemma 4, a lower bound for minimum loss of deviating from $\alpha(\omega^t)$ is

$$\min_{\zeta \in \hat{\Psi}(\omega^t)} \{P(\{\zeta\})\} \times 1. \tag{28}$$

By comparing the above two equations we can see that the maximum gain is strictly smaller than the minimum loss. □

The following proposition follows immediately from lemma 5.

Proposition 6. *If α is consistent, then there is a probability measure that assigns positive probability to any possible finite sequence of observations and a bounded utility function such that with respect to that probability measure, α is the unique optimal contingent plan.*

References

Leo Breiman. *Probability*. Society for Industrial and Applied Mathematics, 1992.

Larry G. Epstein and Michel Le Breton. Dynamically consistent beliefs must be Bayesian. *Journal of Economic Theory*, 61:1–22, 1993.

Edward J. Green and In-Uck Park. Bayes contingent plans. *Journal of Economic Behavior and Organization*, 31:225–236, 1996.

Peter J. Hammond. Consequentialist foundations for expected utility. *Theory and Decision*, 25:25–78, 1988.

Leonard Jimmie Savage. *The Foundations of Statistics*. John Wiley, N.Y., 1954.

Paul Weller. Consistent intertemporal decision-making under uncertainty. *The Review of Economic Studies*, 45:263–266, 1978.

E-mail address: fzb5027@psu.edu

E-mail address: eug2@psu.edu