Max Albert

Bayesian Learning and Expectations Formation: Anything Goes

Appeared in D. Corfield and J. Williamson, eds. *Foundations of Bayesianism,* pp. 341-362, Dordrecht etc.: Kluwer, 2001.

1 Great Expectations

When Muth (1961) introduced the rational-expectations hypothesis (REH), his basic idea was that agents form expectations by rationally acquiring and processing information (weak REH). From this, he immediately jumped to a stronger hypothesis (strong REH), which in the following years radically transformed macroeconomic theory and policy. The strong REH is implied by the assumption that agents know the true (statistical) model of their environment. Except for trivial cases, this model comprises a causal model relating endogenous to exogenous variables, and objective probability distributions of the exogenous variables. Agents' expectations are the objective probability distributions of future developments conditional on their current information about past realizations of exogenous and endogenous variables. In terms of the famous Knightian distinction, the strong REH implies risk and not uncertainty. Rationality then requires that agents choose a strategy (i.e., a plan specifying actions for all contingencies) that maximizes the expected utility on the basis of a v. Neumann–Morgenstern (NM) utility function.[1]

However, the optimistic spirit of the "rational expectations revolution" (Begg 1982) has long since evaporated. Theoretical and empirical weaknesses of the strong REH have become apparent, and there seem to be good arguments for going back to the weak REH. Rationally acquiring and processing information without knowing the true statistical model of the environment (íe., under conditions of *uncertainty* rather than risk) means rational learning, which is the domain of the subjectively expected utility (SEU) theory a.k.a. *Bayesianism.*[2]

[1]Cf. Pearl (2000) for a discussion of causality in relation to statistics. Cf. Hacking (1990) on objective vs subjective probabilities. For strict subjectivists, the strong REH makes sense only for a group of agents, where it translates into the Common Prior Assumption (cf. Aumann 1987: 12ff). See also 2.1 below.

[2]For critical discussions of the strong REH cf. Frydman and Phelps (1983) and Pesaran (1989); for an overview of the learning literature cf. Kirman and Salmon (1995).

The simplest version of Bayesian learning requires the agent to proceed from a *subjective* joint probability distribution for all conceivable future observations. This distribution reflects personal *degrees of belief*. The optimal strategy maximizes *subjectively* expected utility. The initial subjective distribution, the so-called prior distribution, is again revised by conditioning on observed events, which yields the so-called posterior distributions. The whole revision process, which is equivalent to the use of Bayes' theorem, is also called "updating the prior" because the posterior at one stage of the process serves as the prior on the next.

Bayesianism's claim to importance rests on the possible use of a two-stage procedure for deriving the prior. The agent first considers several models (or hypotheses or theories; we use these terms interchangeably) since the true model is unknown. Each model leads to different expectations. Then a prior over the set of models is chosen, which leads to a weighted average of the model-specific expectations. Updating the prior implies a shift in the weights of the models. This two-stage procedure connects Bayesianism with scientific procedures, leading to a unified theory of rationality in economics, statistics and practical decision making.[3]

In a Bayesian context, the strong REH is implied by the assumption that the agent's prior is degenerate and assigns probability 1 to the true model (*true beliefs*). Even without true beliefs, however, rational expectations are possible. Suppose that Adam the Agent tries to predict the outcome of once tossing a fair coin, and that he considers two hypotheses, namely, a probability of head equal to 0.25 or 0.75, respectively. Adam has rational expectations if his prior assigns a probability of 0.5 to each of the two hypotheses, implying that he assigns zero probability to the truth. By definition, rational expectations only require that the subjective probability distribution of the observable variables implied by the prior coincides with the corresponding objective probability distribution (Pesaran 1989: 1).

Does Bayesian learning converge to rational expectations? Again, early optimism turned into disappointment. Even the beliefs of an ideal Bayesian learner who does *not* dismiss the true model from the outset and who faces *no* costs of gathering further information are not inevitably bound to converge to rational expectations (Blume and Easley 1995: 16-20).

It is not clear, however, what lesson, if any, should be drawn from the possible failure of convergence. A non-convergent learning process is not necessarily an in-

[3]The classical exposition of Bayesianism is Savage (1954); see also Kiefer and Nyarko (1995) for a summary and defense emphasizing learning and expectations formation. The approach of Anscombe and Aumann (1963) is appropriate if beliefs refer to objective probability distributions. For a single agent or under the Common Prior Assumption (see 2.1 below), weak and strong REH are formally identical, because the models can be treated like unobservable exogenous variables with a given known distribution (namely, the prior).

dication of misguided decision making (Kiefer and Nyarko 1995). The justification of the Bayesian approach lies not in the convergence properties of Bayesian learning but in the appeal of certain axioms for preferences on the set of strategies. These axioms ensure the existence of an NM utility function and a prior such that the SEU of the strategies reflects the preference ordering.[4]

Therefore, much of the discussion has centered on the Bayesian system of axioms. Nevertheless, we postpone any comments on the axioms until the very end of the paper. Instead, we focus on the fact that Bayesianism is *empty.* As a positive theory, it implies no "operationally meaningful theorems" (OMTs), i.e., consequences that could potentially be refuted by observations.[5] Any behavior can be rationalized on the basis of some prior, even if the NM utility function is given.

For this reason, Bayesianism is also empty as a normative theory. Assume that Mike the Manager asks Betty the Bayesian for advice. Betty cannot take Mike's current beliefs for granted because it is an open question which beliefs are rational given Mike's previous experiences. So it is natural for Mike to ask whether there are any OMTs: Given what I know about the past, and given my NM utility function, is there anything a rational person would *not* do? Is there a sequence of future choices, including reactions to new information, that can be classified as *irrational?* Since the answer is no, Bayesianism is empty as a normative theory. To a Bayesian, *there is no such thing as irrational behavior.*

Although several such rationalizability results have appeared in recent years, they seem to be not widely known, and their implications are not yet fully appreciated.[6] The result discussed in the present paper assumes that an agent considers a simple chaotic process as explanation for observed phenomena. Such explanations are actually considered in economics and elsewhere; they cannot be dismissed as unreasonable. Chaotic systems have the typical features of very rich sets of hypotheses. Since perfectly rational agents by definition consider such very rich sets, they are necessarily, i.e., *independently from the actual complexity of their environment,* in

[4]Preferences over strategies entail preferences over objective-probability distributions of outcomes (expressed by the NM utility function) and beliefs (expressed by the prior). It is fundamental to Bayesianism that preferences in the narrower sense and beliefs are separable (Binmore 1993: 207, Aumann 1987: 13 n. 13). This implies that an agent can adopt an NM utility function independently from her beliefs, or beliefs independently of her NM utility function. On the former case, see also Binmore (1993: 207) on Savage and "massaging the priors". The latter case is illustrated by Savage's own use of the sure-thing principle as a device for (implicitly) adjusting his evalutation of NM utilities in the Allais Paradox (where probabilities are given); cf. Pope (1991).

[5]Samuelson's (1947: 3) phrase, rather than "empirical content", is used to remind readers of the present paper's close relation to Samuelson's work on revealed preference (see 2.3 below).

[6]The present paper is based on Albert (1996, 1999). A slightly different result is contained in Nyarko (1997), who refers to an unpublished 1992 paper of J. S. Jordan for yet another version.

the situation of a person trying to predict a chaotic system. This implies that their expectations are completely arbitrary. Muth's (1961) conjecture that the weak REH provides a solution to the problem of expectations formation is thus refuted, at least if the weak REH is identified with Bayesianism, as it is usually done in economics.

Section 2 reviews the literature and discusses some *prima facie* arguments against the view that Bayesinism is empty. Section 3 introduces an abstract decision problem, section 4 a set of hypotheses based on a simple chaotic system. Section 5 shows that this set can be used to rationalize any strategy. Section 6 concludes with a consideration of arguments against the position, taken in the present paper, that the emptiness of Bayesianism is a serious flaw.

2 A Folk Theorem

Presumably, practitioners tend to believe that there are objectively wrong decisions or mistakes and that decision theory provides the means to avoid them.[7] Theoreticians think differently. It is the folk theorem of decision theory that the notion of rationality employed in economics is "weak".[8] As its counterpart in game theory, the theorem is of unknown origin and implies that (almost) anything goes. Until quite recently, there has been no general proof, but proofs for finite cases are trivial.

There is a difference, however, between the claim that Bayesian rationality is "weak" and that it is completely empty. We therefore discuss three *prima facie* objections to the latter view. 1. It is sometimes suggested that prior beliefs of rational agents are not completely arbitrary. 2. Bayesians have always argued that their definition of rationality implies the rejection of certain other decision rules like the maximin rule; conflict, however, presupposes content. 3. Bayesianism encom-

[7]See, e.g., Bernstein (1996: 336). Goldman (1999: 76) also seems to believe that the so-called Dutch Book argument demonstrates that Bayesianism protects against unnecessary losses. However, the argument assumes a situation without any uncertainty concerning gains or losses and, therefore, completely misses the point when used as a defense of Bayesianism.

[8]Hahn (1996: 186) writes that rationality "buys only a small bit of an answer" in an intertemporal context since it has to be supplemented by a theory on agents' beliefs. Blume and Easley (1995: 26) conclude that the content of Bayesian rationality mostly derives from restrictions on the set of beliefs. Bicchieri (1993: 14, esp. n. 9) restricts the predictive usefulness of Bayesian rationality to stable environments and choice situations familiar to the agent, and mentions convergence problems in the case of complicated priors. Arrow (1990: 29) writes that the rationality hypothesis by itself is "weak" and that its force derives from supplementary hypotheses. By varying utility functions for given beliefs, Ledyard (1986) demonstrates that Bayesianism is empty for a quite general game-theoretic setting. However, he is still convinced of its value as a normative theory (Ledyard 1986: 60, 80f). Bray (1983: 123f) quotes Lucas to the effect that Bayesianism "in many applications" has "little empirical content" but defends it on account of its convergence properties.

passes the theory of demand, which is known to imply OMTs, the so-called axioms of revealed preference.

2.1 Rational Priors

Some Bayesians defend the view that, even though there are no restrictions on priors, all rational agents should hold the same subjective probabilities if they have been exposed to the same experience (e.g., Aumann 1987: 7, 13f). This view is known as the Common Prior Assumption (CPA). Aumann (1987) refers to Savage in this context (without giving a reference) and conjectures that Savage would have accepted the CPA. I disagree (cf. Savage 1962: 11, 13, 14). However, Savage was convinced that *in practice* experience often leads to convergence of opinion. But this is not a starting point for Bayesianism; it is a fact in need of explanation. For convergence, one needs priors that are not too different. The CPA just begs the question in assuming identical priors.

The CPA makes sense only if there exist canonical or rational priors before any experience. This leads to the classical problem of whether there is an acceptable "principle of insufficient reason" determining probabilities before experience. This idea, going back to Laplace, has been criticized by many authors (cf. Leamer 1978: 11, 22-39, 61-63, 111, 114; Howson and Urbach 1989: 45-48, 285, 289; Earman 1992: 14-17, 138-141). It had been revived by Keynes and others in the form of a theory of "logical" probabilities, i.e., uniquely determined *a priori* probabilities of the logically possible hypotheses. One of the arguments in favor of Bayesianism has been the discovery that such probabilities do not exist.[9] It seems not to be a promising way of further development to revive this idea again. As the history of the subject presents itself, the burden of proof that there is an acceptable "principle of insufficient reason" rests with those in favor of the CPA.

2.2 The Dominance Principle

Clashes between Bayesianism on the one hand and decision rules for behavior under uncertainty like the maximin rule on the other hand are due to the fact that the latter violate what we will call the dominance principle.[10] We can use the NM

[9]This appears already to have been a conjecture of Ramsey, the earliest of the modern Bayesians, who made this argument against Keynes; cf. Hacking (1990: 165, 170). The Keynesian program was taken up later by Carnap; it was intended to provide one of the cornerstones of logical positivism. There is a widespread agreement today that this program foundered in just the way Ramsey conjectured: there are no logical probabilities; cf. Howson and Urbach (1989: 48-56).

[10]This clash has most often been stressed in connection with statistical decision theory, see, e.g., Lindley (1972: 13-15).

utility function to define a set of *strictly dominated strategies.* A strategy A is *strictly dominated* if and only if for every choice of prior probabilities, there always exists at least one strategy with a higher SEU than A. Bayesianism implies one restriction on behavior for a given NM utility function, namely, that no strictly dominated strategy is chosen. Let us call this restriction the dominance principle.

We illustrate this principle for a case of three strategies leading to different consequences in two mutually exclusive and jointly exhaustive states (figure 1). Given an NM utility function, the three strategies can be represented by their utilities in a two-dimensional coordinate system. Bayesian analysis implies the linearity of the indifference loci in this diagram; the maximin criterion would lead to L-shaped indifference loci. Therefore, the latter criterion allows for choices that are ruled out by Bayesian analysis.

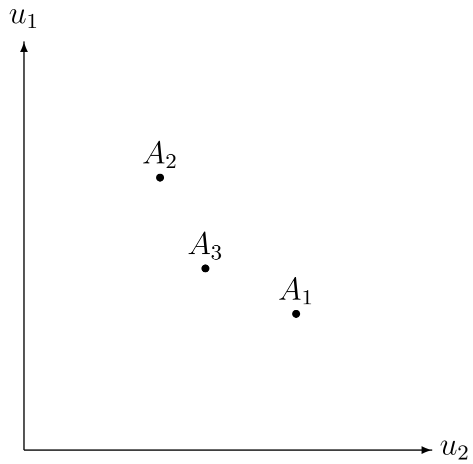

Figure 1: A strictly dominated strategy. The axes measure the NM utilities of three strategies in the case of two mutually exclusive and jointly exhaustive states. The dominated strategy A_3 is selected by the maximin criterion.

The set of strictly dominated strategies is by definition independent of beliefs. Nevertheless, the dominance principle yields no OMTs. Identifying strictly dominated strategies already requires at least some knowledge about one's environment. If the NM utility function is given, it is always possible to imagine, for each strategy A, a state s such that A yields a higher utility than any other strategy if s holds. The assumption that a strategy is strictly dominated means that logical possibilities are excluded. If Bayesianism is empty as a theory of learning, the dominance principle yields neither predictions nor advice, because before anything has been learned, logical possibilities *cannot* be excluded.

2.3 The Revealed-Preference Approach

In economics, it is usually taken for granted that economic theory only provides assumptions about the *general structure* of agents' preferences, and that the details necessary for making predictions have to be gathered by observation. According to this view, economic theory provides a framework that allows us to use empirical observations to make predictions, much as Newton's theory allows us to use empirical observations to determine the masses of the planets, which in turn can be used to predict the planets' movements.

An important question in this context is which kind of observations can be used. The economic tradition allows only the use of observations about actual choices. This is consistent with the idea that—presumably as a result of selective pressure—agents act *as if* they were perfectly rational, or at least *as if* deviations from perfectly rational behavior were unsystematic. The as-if approach views rationality as a feature of behavior, not as a feature of the process of deliberation.[11]

This attitude is the basis of the revealed-preference (RP) approach to the theory of demand (cf., e.g.,Varian 1992: 133). The approach focuses on OMTs—usually called *axioms* of RP since they exhaust the theory—implying that information about past choices can be used to rule out certain future choices if preferences have a certain structure.

The theory of demand implicitly assumes that observer and agent share true beliefs in a single deterministic model. Bayesianism can only profit from the results of this theory if the observer can check whether the agent holds appropriate beliefs. According to the economic tradition, such a check must be based on observations of actual choices and nothing else. Hence, an extension of the RP approach to the case of uncertainty is required. Such an extension would provide OMTs demonstrating how, assuming the axioms of Savage (1954) or Anscombe and Aumann (1963), information about past choices can be used to rule out certain future choices.

There are some extensions of the RP approach, but none that cover Bayesian learning.[12] The present paper addresses this issue in a very general way. As a concession to Bayesianism, knowledge of the NM utility function summarizing an

[11] Cf. Simon's (1976) characterization of economic rationality as "substantive" rather than "procedural", which seems to be meant as a clarification of Friedman's (1953) as-if approach.

[12] Border (1992) develops an RP approach to choice among lotteries with monetary rewards. The observer knows only that more money is preferred to less. If observer and observed agree on all (objective) probabilities, any choice behavior that is not stochastically dominated can be rationalized by postulating a suitable utility function. The RP approach of Green and Osband (1991) is based on assumptions that deviate from Savage's (1954) framework in several ways. A direct comparison of results is therefore difficult. Kim (1992) considers choice under uncertainty but excludes learning, i.e., conditionalization on past observations.

agent's behavior under certainty and risk is granted. This leaves only beliefs as missing determinants of behavior. An RP approach to Bayesian learning requires OMTs implying that, given the NM utility function, information about past choices can be used to rule out certain future choices.

3 Questions

The technical aspects of our problem, and the way to a solution, can be explained with the help of a simple decision problem; generalizations are trivial (see 5.2 below).

Adam & Eve and the Money Spinner. Adam the Agent owns a mysterious black box connected with a screen and a keyboard. The screen displays either 0 or 1; in fixed intervals the screen goes black and then again shows one of the two digits. The tth observation is denoted by x_t. At every point in time $t = 0, 1, \ldots, \infty$, Adam places his bets on the next digit by typing in a number $y_t \in \{0, 1, ..., N\}$. Doing nothing implies $y_t = 0$. So the sequence of events is as follows. At $t = 0$, Adam chooses y_0. At $t = 1$, x_1 appears and Adam chooses y_1. And so on. At time t, Adam chooses y_t. If $x_{t+1} = 0$, the black box produces y_t perfect one-dollar notes. If $x_{t+1} = 1$, it produces $N - y_t$ equally perfect one-dollar notes. Adam cares only for the money he receives; his NM utility function $\nu\colon [0, N] \mapsto \mathbb{R}$ at each point in time is increasing, strictly concave, and finite. Adam's information at time t encompasses all the facts just explained and the history of digits and choices up to time t.

Eve the Economist observes Adam. Her information at time t coincides with Adam's; specifically, she knows his NM utility function. We simplify the problem by assuming that both Adam and Eve know (i.e., rightly believe) that Adam's choices have no influence on the sequence of digits appearing on the screen. ■

Let X be the set of all finite sequences or histories of observations (0s and 1s). These sequences are of varying length; $\ell(\mathbf{x})$ is the length of $\mathbf{x} \in X$. Similarly, let Y be the set of all finite sequences of choices (natural numbers from 0 to N), with $\ell(\mathbf{y}) = t$ as the length of $\mathbf{y} \in Y$. In both cases, we include the sequences of zero length ("vacuous" histories). The set XY denotes all pairs $(\mathbf{x}, \mathbf{y})$ from $X \times Y$ with $\ell(\mathbf{x}) = \ell(\mathbf{y})$.

The following three questions define the problem we are interested in by recourse to the situation of Adam and Eve.

Question 1 *Can Eve exclude some histories $(\mathbf{x}, \mathbf{y}) \in XY$ as inconsistent with the assumption that Adam is a perfectly rational Bayesian agent?*

Question 2 *Can Eve, on the basis of finitely many observations $\mathbf{x} \in X$, give good advice to Adam from a Bayesian point of view?*

Question 3 *Assume that Eve is a Bayesian-minded economist who has observed a sequence of digits and choices* $(\mathbf{x}, \mathbf{y}) \in XY$*. Is there any restriction on Eve's beliefs concerning Adam's future behavior resulting from the hypothesis that Adam is a perfectly rational Bayesian agent?*

Question 1 concerns Bayesianism as a positive theory that should yield predictions of Adam's behavior. Question 2 concerns Bayesianism as a normative theory that could be used by Eve to advise Adam. Question 3 concerns Bayesianism as a methodology used by Eve to analyze Adam's behavior.

Of course, all three questions are strongly interrelated. If Adam can rationalize any choice of strategy, question 1 must be answered in the negative. The same goes for question 2. If any strategy can be rationalized, there is nothing a Bayesian advisor can say except "Do what you want".

The answer to question 3 is slightly more involved. There is a difference between questions 1 and 3. A negative answer to question 1 implies that no OMTs exist. But a Bayesian could still claim that Bayesianism as a methodology allows one to conclude that certain sequences of digits and choices become very improbable if Adam is rational. If that were possible, it would provide a Bayesian argument against the requirement that positive theories should provide OMTs.

Let us shortly summarize the Bayesian analysis of Adam's problem. First of all, Adam should choose a set $\mathcal{H}$ of mutually exclusive hypotheses, each of which implies objective probabilities for all potential future observations. Then, he should choose a subjective probability measure or prior μ on (a σ–algebra of subsets of) $\mathcal{H}$. The prior μ is chosen such that it generates a preference ordering on the set of all strategies. Hence, we know that, for every history $\mathbf{x} \in X$, the pair $(\mathcal{H}, \mu)$ implies conditional probabilities $P_\mu\big(x_{\ell(\mathbf{x})+1} = i \,|\, \mathcal{H} \wedge \mathbf{x}\big)$, $i = 0, 1$.[13]

Adam knows that there is no influence from his choice at one point in time to consequences at other points in time. The only connection between choices is learning. Given a "forecast function" (Nyarko 1997: 181) $\rho \colon X \mapsto [0, 1]$ assigning the probability $p = \rho(\mathbf{x})$ to the event $x_{\ell(\mathbf{x})+1} = 0$ (the next digit is 0), the action at $t = \ell(\mathbf{x})$ can be considered separately from other actions. Adam maximizes his SEU, solving the problem

$$\max_y \{p\nu(y) + (1-p)\nu(N-y) \colon p = \rho(\mathbf{x}),\ y \in \{0, \ldots, N\}\}. \tag{1}$$

[13] The symbol $\wedge$ in $\mathcal{H} \wedge \mathbf{x}$ denotes the conjunction. Read as statements, $\mathcal{H}$ is a (possibly uncountable) disjunction of hypotheses, and $\mathbf{x}$ is equivalent to the conjunction of statements "At time s, x_s is observed", $s = 1, \ldots, \ell(\mathbf{x})$. If $\ell(\mathbf{x}) = 0$, $\mathcal{H} \wedge \mathbf{x}$ is of course equivalent to $\mathcal{H}$. If Adam observes a sequence $\mathbf{x}$ with subjective probability 0, the conditional probabilities $P_\mu\big(x_{\ell(\mathbf{x})+1} = i \,|\, \mathcal{H} \wedge \mathbf{x}\big)$ are not defined. He is then free to choose a new prior distribution, which does not improve the chances of predicting his actions. However, we can exclude this case (see below).

His optimal strategy, then, is described by the function $y\colon X \mapsto \{0, \ldots, N\}$ assigning the utility-maximizing choice to each conceivable history:

$$y(\mathbf{x}) \stackrel{\text{def}}{=} \arg\max_{y} \{p\nu(y) + (1-p)\nu(N-y)\colon p = \rho(\mathbf{x}),\ \ y \in \{0, \ldots, N\}\} \qquad (2)$$

His actual choice $y_{\ell(\mathbf{x})}$ of course depends on the actual history $\mathbf{x}$.

Bayesian rationality requires that Adam's forecast function reflects some prior:

$$\exists(\mathcal{H}, \mu)\forall \mathbf{x} \in X \big[\rho(\mathbf{x}) = P_\mu\big(x_{\ell(\mathbf{x})+1} = 0 \,|\, \mathcal{H} \wedge \mathbf{x}\big)\big] \qquad (3)$$

If the strategy $y(\mathbf{x})$ is optimal given the set of hypotheses $\mathcal{H}$ and the attached prior μ, the pair $(\mathcal{H}, \mu)$ is a rationalization of $y(\mathbf{x})$.

The analysis is simplified by the fact that a strategy can be rationalized, if at all, without assuming any of the conditional probabilities to be 0, because Adam chooses between discrete values. Given discreteness, any decision that is optimal if some event has zero probability will also be optimal if the probability of the respective event is small enough. Of course, all policy variables in real-world decision problems are discrete since the precision of measurement is always finite. Moreover, each possible choice y in (1) is optimal for some values of p. If, therefore, we can find a rationalization for arbitrary forecast functions $\rho\colon X \mapsto (0, 1)$, where $(0, 1)$ is the open unit interval, we can rationalize any strategy $y\colon X \mapsto \{0, \ldots, N\}$.

4 The Chaotic Clock

The problem in finding rationalizations is that a set of hypotheses might not be rich enough to provide a rationalization for a given strategy. However, there are very trivial sets of hypotheses that are always rich enough.

4.1 The Basic Mechanism

Assume that the evolution of the inner states of Adam's black box follows a deterministic process depending on a starting point. The law of the deterministic process is the baker-map dynamics, which can be graphically illustrated as the output of a chaotic clock (figure 2).

There is one pointer that can point to all real numbers in the intervall $I = [0, 1)$, where the vertically upward position is zero and the vertically downward position is $\frac{1}{2}$. Initially, the pointer deviates by an angle $\omega = 2\theta\pi$ from the vertically upward position, thus pointing at the real number θ. At $t = 1, 2, \ldots, \infty$, the pointer moves by doubling the angle ω. If the pointer comes to rest in the first half of the dial and

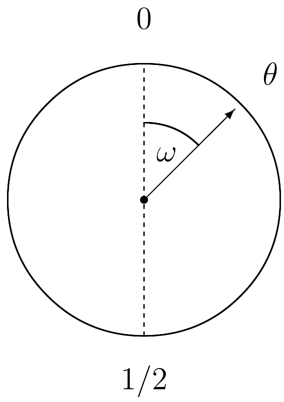

Figure 2: A chaotic clock. At each point in time, the angle ω is doubled. When the pointer is in the first (second) half of the dial, the digit 0 (1) appears on a screen. The resulting sequence of digits is described by the baker-map dynamics.

points at a number in $\left[0, \frac{1}{2}\right)$, the screen of Adam's black box shows 0; otherwise it shows 1.

According to the chaotic-clock hypothesis, the inner states of the black box at $t = 0, 1, \ldots, \infty$ are described by a real variable, the pointer position $z_t \in I$. The inner states evolve deterministically, but it can only be observed whether the pointer position is $z_t \in \left[0, \frac{1}{2}\right)$ or $z_t \in \left[\frac{1}{2}, 1\right)$. These two states result in $x_t = 0$ and $x_t = 1$, respectively. The deterministic law by itself does not allow for a prediction of future observations; an assumption concerning the starting point $z_1 = \theta$ is also necessary. Thus, there is a set of hypotheses, one for each starting point $\theta \in I$. The corresponding dynamical system is the baker-map dynamics:[14]

$$\begin{array}{lrcl} \text{(a)} & x_t &=& g(z_t) \overset{\text{def}}{=} 2z_t \text{ div } 1 \\ \text{(b)} & z_{t+1} &=& h(z_t) \overset{\text{def}}{=} 2z_t \mod 1 \\ \text{(c)} & z_1 &=& \theta \end{array} \tag{4}$$

Note that the chaotic clock cannot produce an unbroken infinite sequence of 1s. If the pointer is in the second half of the dial (thus generating a 1 on the screen), doubling the angle ω moves the pointer beyond 0, thus leading to a *smaller* value

[14] "div" denotes integer division; "mod" denotes the indivisible rest of the integer division, i.e., $x \text{ mod } n \overset{\text{def}}{=} x - (x \text{ div } n)$. On the baker-map dynamics, see Ford (1983), Devaney (1989: 18 example 3.4, 39, 52) and Schuster (1988: 107f). The graphical illustration is due to Davies (1987: ch. 4).

of ω. As long as the pointer comes to rest in the second half of the dial, ω falls at every tick of the clock with increasing rates, until, after a finite number of ticks, the pointer is in the first half of the dial, which implies that the screen shows 0.

4.2 Falsification Dynamics

Assume that Adam believes that his black box contains a chaotic clock. In order to analyze the consequences of uncertainty concerning θ, Adam has to know how x_t develops for a given θ. This is very simple in principle. Every $\theta \in I$ can be expressed as

$$\theta = \sum_{n=1}^{\infty} \frac{\theta_n}{2^n} = 0.\theta_1\theta_2\theta_3 \cdots \theta_t \cdots \tag{5}$$

where θ_n is 0 or 1 (dyadic development). In order to enforce uniqueness of the representation (5), we require infinitely many 0s on the right-hand side. Thus, $\frac{1}{2}$ should be represented as $0.1\overline{0} = 0.1$ and not as $0.0\overline{1}$ (where the bar denotes infinite repetition).[15]

The sequence generated by the chaotic clock is just the dyadic development of the starting point, i.e., x_t is equal to θ_t in (5). This follows from two mathematical facts. (i) We have $\theta \geq \frac{1}{2}$ if and only if $\theta_1 = 1$. (ii) Doubling θ shifts the point in the dyadic development one step to the right; the modulo function sets the digit before the point to 0, thus ensuring that the new value is again smaller than 1.

The first digit θ_1 in (5) determines whether θ is smaller than $\frac{1}{2}$ or not. Hence, when Adam observes the first digit $x_1 = \theta_1$, he finds out whether the starting point θ is in the first half ($x_1 = 0$) or in the second half ($x_1 = 1$) of the dial. The second digit places the starting point into one of the four quarters. For example, if $x_1 = 1$ and $x_2 = 0$, the starting point must be in the third quarter $\left[\frac{2}{4}, \frac{3}{4}\right)$. And so on. The sequence of digits corresponds to a sequence of bisections of the set I of potential starting points; with each further digit, the location of the starting point is narrowed down to the upper or lower half of the remaining interval on the dial. At time t, Adam has made t observations revealing one of the 2^t *basic intervals*

$$I_t(m) \stackrel{\text{def}}{=} \left[\frac{m}{2^t}, \frac{m+1}{2^t}\right), \; m \in \left\{0, \ldots, 2^t - 1\right\} \tag{6}$$

as the location of the starting point θ. In infinite time, these bisection steps converge to θ.

In other words, each observation falsifies half of the remaining hypotheses concerning the starting point θ. Hence, Adam's beliefs converge to the truth if he is

[15] On coin tossing and dyadic development see also Brémaud (1988: 28-31).

right in assuming that the black box contains a chaotic clock. This does not mean, however, that his chances of predicting the future improve over time. Obviously, knowing the first t digits of θ's dyadic development is no help in predicting the next digits.

4.3 Souping up the Clock

The chaotic clock places one restriction on Adam's system of beliefs: it is impossible that the number of 0s is finite. In order to get rid of this restriction, we assume that the angle between the pointer and the zero position is not doubled as before but quadrupled:

$$\begin{array}{lrcl} \text{(a)} & x_t & = & g(z_t) \\ \text{(b)} & z_{t+1} & = & h(2z_t) \\ \text{(c)} & z_1 & = & \theta \end{array} \tag{7}$$

The functions g, h are defined as in (4). For any starting point θ with dyadic development (5), the original system (4) yields the observations $x_t = \theta_t$. The modified system (7) yields $x_t = \theta_{2t-1}$, i.e., every second digit of the the starting point's dyadic development is irrelevant. Therefore, starting points like $\theta = 0.\overline{10}$ or $\theta = 0.\overline{1011}$ generate an unbroken infinite number of 1s although they feature infinitely many 0s as required.

Almost the same analysis as before applies to the modified chaotic clock. Since every second digit of the starting point's dyadic development never appears on the screen, an infinite number of starting points lead to the same sequence of observations. The first t observations determine 2^{t-1} basic intervals of length $\frac{1}{2^{4t-2}}$, one of which contains the starting point. Each basic interval is as rich as the unit interval. It still contains points that can produce any kind of extension to the observed sequence. In other words, the set of hypotheses represented by the set of starting points is so rich that, for any possible future, there are (actually: infinitely many) hypotheses consistent with the past and predicting this future. *No future developments can be excluded on the basis of past observations.*

Adam's beliefs still converge to the truth but not to the full truth: the observations reveal a subset of possible starting points containing the true one. This set corresponds to a subset of observationally equivalent hypotheses, one of which is true. As before, convergence to the truth presupposes that the black box actually contains a modified chaotic clock. If this is not the case, however, the analysis of the learning process is unaffected since no sequence of observations, not even an infinite one, is inconsistent with the assumed mechanism.

However, Adam does not *know* whether the black box actually contains a chaotic clock. One may ask, then, why he should ever consider such a crazy mechanism.

There are several reasons.

First of all, Adam is a perfectly rational being and, therefore, logically omniscient (cf. Earman 1992: 121f). He is aware not only of the chaotic-clock hypotheses but of many more when he considers the question of choosing a prior. Therefore, the question should rather be why not more hypotheses are considered.

Secondly, if the truth is deterministic, it is observationally equivalent to (7) (with a specific starting point). Thus, (7) can be viewed as a "reduced form", as econometricians would say, of any deterministic theory (plus initial conditions) that is capable of explaining an infinite sequence of the events in question. In this sense, the chaotic clock represents all deterministic theories, and Adam's problem does not become significantly greater when we add further hypotheses to the set of chaotic-clock hypotheses.[16] Hence, considering the chaotic clock is actually a simplification (although an insignificant one) of Adam's problem.

Last but not least, the chaotic-clock model fulfills all the formal requirements of a scientific theory. It assumes a simple mechanism governed by a law of motion that produces different results according to the initial position of the mechanism. While not even the Swiss could actually produce the chaotic clock, processes that lead to chaotic dynamics are not rare, and imperfect observability can produce the kind of irregular behavior characteristic of the chaotic clock. Moreover, although the chaotic-clock model assumes a continuum of starting points (corresponding to more hypotheses than we could ever consider explicitly), it is less complicated than models encountered in physics or economics. It would be difficult to find any acceptable formal requirement that excludes the chaotic clock from consideration.[17]

Let us call "empirically adequate" those hypotheses that are observationally equivalent to the truth. Subsequently, convergence to the truth is taken to mean that, in the limit, the subjective probability of the set of empirically adequate hypotheses is 1. If the truth is deterministic, then, convergence to the truth is ensured.

5 Answers

5.1 Adam's Problem

We turn to the dynamics of probabilistic beliefs generated by the modified chaotic-clock hypotheses described by (7). The hypotheses form a set $\mathcal{H}^*$, where $H_\theta \in \mathcal{H}^*$,

[16] Even adding probabilistic hypotheses would make no difference, as will become obvious below.

[17] The chaotic clock poses a generalized version of Goodman's (1955) "new riddle of induction". The set of hypotheses considered by Goodman is countable and, therefore, too small to lead to the problems discussed in the present paper. Using the chaotic clock for presenting the problem of induction has the advantage that no "gruesome" predicates appear.

$\theta \in I$ denotes the single hypothesis that the starting point of the modified chaotic clock is θ. Since these hypotheses are deterministic, the probabilities $P(x_t = 0 \,|\, H_\theta)$ are either 0 or 1, i.e., H_θ yields certain or point predictions. Uncertainty enters via the uncertainty concerning $\theta \in I$. As a Bayesian, Adam chooses a subjective probability measure on (a σ–algebra of subsets of) $\mathcal{H}^*$, and since each hypothesis in $\mathcal{H}^*$ is represented by a $\theta \in I$, we can consider instead a subjective probability measure on (a σ–algebra of subsets of) I.

Adam needs well-defined conditional probabilities $P_\mu\big(x_{\ell(\mathbf{x})+1} = i \,|\, \mathcal{H}^* \wedge \mathbf{x}\big)$ for any potential sequence of observation $\mathbf{x}$. Hence, he must include all the basic intervals $I_t(m)$ in his σ–algebra of subsets of I. We can therefore consider the σ–algebra generated by the basic intervals (which is just the σ–algebra of the Borel sets). Any probability measure μ on this σ–algebra then determines the conditional probabilities Adam needs to solve his decision problem. Henceforth, we simply speak of a probability measure μ on I or, equivalently, on $\mathcal{H}^*$.

We have already seen in section 3 that the central question is if we can find a rationalization $(\mathcal{H}^*, \mu)$ for arbitrary forecast functions $\rho\colon X \mapsto (0,1)$. This question is answered by the following theorem (Albert 1999: theorem 1).

Theorem (Anything Goes) *Let $\mathcal{H}^*$ be the set of modified chaotic-clock hypotheses. Consider an arbitrary forecast function $\rho\colon X \mapsto (0,1)$. Then, there exist infinitely many probability measures μ on $\mathcal{H}^*$ such that the rationality condition $\rho(\mathbf{x}) = P_\mu\big(x_{\ell(\mathbf{x})+1} = 0 \,|\, \mathcal{H}^* \wedge \mathbf{x}\big)$ holds for all $\mathbf{x} \in X$.*

In interpreting the theorem, we have to remember that Adam as a perfectly rational person is always aware of the implications of all the assumptions he is considering. When choosing the prior μ on $\mathcal{H}^*$, he is aware of the implicit assignments of numerical values to the conditional probabilities. The theorem says that, instead of choosing a prior on $\mathcal{H}^*$, Adam may as well choose arbitrary conditional probabilities.[18] And since these probabilities can generate any contingent choices whatsoever, it is immaterial for Adam whether he asks himself "What should I do?" or "What prior should I choose?". The Bayesian apparatus provides no restrictions and therefore no help in making this choice.

Let us consider two simple cases. Adam might, for instance, choose a constant forecast function $\rho(\mathbf{x}) = \frac{1}{2}$ for all $\mathbf{x} \in X$. This is rationalized by the uniform prior with density $f(\theta) = 1$ on I. This choice implies that, independently from past observations $\mathbf{x}$, the probability of 0 and 1 is always $\frac{1}{2}$. Any other constant forecast function leads to a misleadingly complicated prior distribution that has no density.[19]

Generalizing slightly, Adam might set $\rho(\mathbf{x}) = p[\ell(\mathbf{x})]$ with an arbitrary function

[18]Thus, including probabilistic hypotheses in addition to $\mathcal{H}^*$ would not change the results.

[19]Cf. Brémaud (1988: 29) who, however, discusses only the uniform distribution.

$p\colon \mathbb{N} \mapsto (0,1)$, thus fixing his decision at time t independently from and prior to any observations. Adam poses as a Bayesian learner but is actually completely "dogmatic" (and completely unpredictable) in the sense of ignoring any experience.

A Bayesian dogmatic is already unpredictable. A run-of-the-mill Bayesian, who allows experience to influence his behavior whenever it suits him, is all the more unpredictable because he has more options. This answers question 1. Bayesian rationality is empty as a positive theory. Eve cannot exclude any behavior of Adam on the basis of the hypothesis that Adam is rational in the Bayesian sense. Nor can Eve give any advice to Adam, even if she knows his NM utility function, since no strategy is irrational, whatever observations have been made. This answers question 2: Bayesianism is empty as a normative theory.

In relation to the tent-map dynamics (which is mathematically equivalent to (4)), Blume and Easley (1995: 19f, 36f) show that convergence to the truth need not mean that predictions improve: if the prior is continuous at the starting point, the posterior distribution converges to the uniform distribution, implying a probability of 0.5 for observing 0. Hence, convergence to the truth does not imply convergence to rational expectations.[20]

This is an interesting point but not the one we are making. Blume & Easley's point provides no argument against Bayesianism (and is not meant to do so), because no procedure can improve on Bayesian learning in these cases. We can get more mileage out of the machinery of chaotic dynamics.

In Blume and Easley's analysis, the *true* law of the process generating the 0s and 1s is chaotic, and this law is *known* to the agent. In our analysis, the true law governing the sequence of 0s and 1s is unimportant. The anything-goes theorem neither refers to some true law nor places restrictions on the sequence of events observed by the decision maker. The problem is not the complexity of the environment but the complexity of a large set of hypotheses. The chaos is in the decision maker's head.

The assumptions of the present analysis are actually quite favorable to Bayesianism since convergence to the truth (in the empirical-adequacy sense) is ensured. Allowing for probabilistic hypotheses would open up the possibility of non-convergence to the truth.[21] As Gillies (2001) shows, a Bayesian learner might then never discover that he had dismissed the truth from the outset, which provides a strong argument in favor of seriously considering a large set of hypotheses.[22]

[20]Hence, "merger of opinions" for different persons (cf. Earman 1992: ch. 6) has nothing to do with convergence of expectations (opinions concerning the future).

[21]Unfortunately, non-convergence cannot be quite as dramatic as suggested by theorem 2 in Albert (1999), which is incorrect. Under the conditions of theorem 2, convergence to rational expectations is ensured although the probability of convergence to the truth may be 0 (as is typically the case if, e.g., both hypotheses assign a probability of 0.5 to observing 0 in the limit).

[22]A hypothesis is dismissed from the outset iff it is not in the support of the prior. The support

5.2 Eve's Problem

We turn to Eve's prediction problem, i.e., to the problem of an economist trying to predict the behavior of rational agents and using Bayesianism as a methodology. If Eve is a Bayesian herself, she is not overly impressed by the fact that Bayesianism yields no OMTs concerning Adam's behavior. The Bayesian methodology allows a continuum of beliefs between the two categories "ruled out" and "possible" that are considered when we speak of OMTs. However, this is not going to help Eve. Her hypothesis that Adam is rational provides no restrictions concerning Adam's behavior. Therefore, she is in no better position to predict Adam than to predict the digits generated by the black box.

For a formal proof, we have to generalize the previous results to larger spaces of observables. This presents no difficulties.

Eve observes the digits on the screen and Adam's choices. Moreover, she might observe other things, like Adam's facial expression or his pattern of consumption, that are or are not related to Adam's behavior. Sticking to our premise that, realistically, all observable variables can only range over a finite set of values, we assume that Eve's observable universe can be "digitalized": each state can be described by a binary string of 0s and 1s with maximum length $n \geq 2$. If the number of different states is between 2^{n-1} and 2^n for some n, several strings describe the same state.

Again, we look for a set of hypotheses sufficiently rich to allow Eve to rationalize any forecast. Such a set is again provided by the (modified) chaotic clock. Assume that Eve considers Adam and the money-spinner as a big black box that displays one of 2^n combination of observables at each point in time. The combinations are determined by a chaotic clock that makes n angle-quadrupling ticks at each point in time; the resulting string s_t of n digits is then revealed as a solid block instead of a succession of digits. Formally, we can leave the chaotic clock as it is; we just have to assume that, at each $t = 1, \ldots, \infty$, s_t is observed instead of just one digit x_t:

$$\begin{array}{llll} \text{(a)} & x_\tau = g(z_\tau) & \text{(b)} & z_{\tau+1} = h(2z_\tau) \\ \text{(c)} & z_1 = \theta & \text{(d)} & s_t = \left\{x_{(t-1)n+1}, \ldots, x_{tn}\right\},\ t = \tau \text{ div } n \end{array} \tag{8}$$

This dynamical system is identical to (7) except for the fact that at each point in time t, the string s_t produced by the last n ticks of the clock is observed. The sequence generated by the system is the dyadic development of the starting point with every second digit removed.

---

is a set with zero-probability complement; sometimes, it is also required that the support's intersection with any open set, if not empty, has positive probability. In our analysis, the support is I since all open sets are measurable and contain basic intervals, which never have probability 0.

The chaotic clock is a "theory of everything" for universes the evolution of which can be described by an infinite sequence of binary string of length n. Since the possibilities of assigning probabilities to sequences are not affected by the fact that these sequences are now revealed in a blockwise fashion, the previous results still apply: according to the anything-goes theorem, arbitrary assignments to these probabilities are possible. This answers question 3. Bayesianism as a methodology is completely useless in predicting rational behavior because there are no OMTs covering this behavior. Eve's expectations concerning Adam's behavior are completely arbitrary.

For example, Eve may decide to view Adam as attaching equal probabilities to two arbitrary hypotheses like "The probability of $x_t = 0$ is $\frac{1}{5}$" and "The probability of $x_t = 0$ is $\frac{2}{5}$" while she herself attaches a probability of $\frac{1}{2}$ to $x_t = 0$. She then can find a prior such that exactly the conditional probabilities implied by this view will hold, no matter what happens.

Moreover, the analysis of Eve's problem shows that the restriction of Adam's problem to the prediction of single digits is immaterial. Everything works as before as long as there is a maximum amount of information he can access at each point in time.

6 Conclusions

Bayesian rationality becomes empty if the decision maker considers a set of hypotheses that is as large as the set described with the help of the chaotic clock. Whatever the actual process generating a sequence of observations, considering a chaotic-clock explanation already implies that any experience can be accommodated without implications for expectations concerning the future. The inclusion of further hypotheses does not add to the complexity of the learning problem.

This is just another version of the problem of induction. Logically, one can never infer the laws governing the world from a finite number of past observation. While many theories may be eliminated over time, it is quite trivial that there always remain enough theories consistent with any kind of future. The so-called pragmatic problem of induction says that learning, if guided by experience and deductive logic alone, yields no restrictions for decision making.[23] The anything-goes theorem shows that Bayesianism, although employing more than just deductive logic, cannot solve the pragmatic problem of induction either.

This result is no surprise once it is clear how many degrees of freedom Bayesianism leaves to the decision maker in setting up the initial beliefs. But it is surprising

[23] Cf. Musgrave (1989: section 4) and Miller (1994: 20-23, 38-45), whose solution rests on the assumption that it is possible to reduce the number of acceptable theories drastically.

that a sufficiently rich set of hypotheses can be introduced in such a simple and compact way. If not much sophistication is needed to experience the problem generated by too many logical possibilities, maximum sophistication or perfect rationality will necessarily lead to this problem.

It is a mathematical fact that any strategy can be rationalized even for a given NM utility function. Several arguments might be raised against the position, taken in the present paper, that this fact speaks strongly against adopting Bayesianism as a positive or normative theory. Specifically, four arguments seem to be important. Their discussion will conclude this paper.

6.1 Supplementing Bayesianism

It is not necessarily alarming if a theory is devoid of empirical content. As long as the theory is not analytic, it might still be possible to supplement it by further hypotheses, creating a larger theory the empirical content of which comes neither from the original theory alone nor from the supplement alone. The same goes, mutatis mutandis, for a normative theory. Even if Bayesianism allows one to rationalize any strategy, there might be supplementary rules that distinguish between good and bad rationalizations.[24]

However, such supplementary rules are equivalent to a principle of insufficient reason. This can easily be shown. Prima facie, there seem to be two different options. Supplementary rules or hypotheses providing content can either a) restrict prior beliefs *on the basis of* experience or b) restrict prior beliefs *without any basis in* experience.

However, option a) is really not different from option b). Assume that there is a rule R that restricts the prior on the basis of experience. Thus, for every set of data E, the rule R selects a prior. This is exactly what Bayesian updating does on the basis of a still earlier prior chosen before E becomes known. We have seen that a decision maker can choose the prior before experience such that, for every set of data E, an arbitrary predetermined posterior results. Thus, whatever the rule R, a prior can be chosen before any experience that is in conformity with the recommendations of R. It follows that R can be replaced by restrictions on the admissible set of priors before any experience; we are in effect left with option b).

A rule determining a prior, or at least restricting the choice of priors, without any basis in experience is a "principle of insufficient reason". As has been argued before, Bayesianism is the product of a history of failures to provide such a principle. Thus, embracing option b) looks not very promising for normative Bayesianism.

[24]Nyarko (1997: 176) makes this point but just provides results concerning the implications of different restrictions placed on priors.

For Bayesianism as a positive theory, of course, it would be irrelevant whether a “principle of insufficient reason” looks reasonable or not as long as the package is successful empirically. However, this kind of success is missing so far, or at least this seems to be the public opinion in economics.

Given the record of Bayesianism as a positive theory of behavior, the advocates of Bayesianism in economics have stressed Bayesianism’s methodological virtues as a systematic theory of behavior as compared with the “adhockery” of bounded-rationality approaches.[25]

As we have seen, Bayesianism does not provide a theory of behavior on its own. Without a theory of priors, the actual hypothesis in each application must needs be chosen in an ad hoc fashion by selecting a prior or an admissible set of priors. Thus, Bayesianism as it stands is not methodologically superior to the bounded-rationality approach. It might be true that we can guess in each application what the beliefs of the agents will be. On the other hand, we might equally well guess which of a set of rules of thumb they are likely to use. The degree of adhockery on both sides seems comparable.

6.2 Bayesianism as Adaptation

Another defense of Bayesianism is the idea that Bayesian perfect rationality might be an idealization anticipating—prematurely, so to speak—the long-run effects of adaptation and training. This means that selective pressures favor Bayesian perfect rationality in the long-run.

Let us compare the hypothetical fate of a perfectly rational with that of a boundedly rational agent. Let us call these hypothetical individuals Priscilla and Brian, respectively. Brian is not logically omniscient; he does not consider a set of hypotheses sufficiently rich to rationalize any behavior. Brian starts with a restricted set of hypotheses, and, for whatever reasons, there are only some priors that appeal to him. Moreover, he is bound to make logical mistakes; thus, even if he tried to maximize his subjectively expected utility, he would often fail to recognize the subjectively optimal actions. On the other hand, Brian might just adopt some rule of thumb for decision making and ignore his own beliefs. Would Brian have any disadvantages as compared with Priscilla? Will the Brians of this world either learn to mimick Priscilla’s cleverness or vanish in the long run?

It seems not, at least from a Bayesian point of view. The difference between

[25]For this argument and the next, cf. Selten’s account of a fictitious discussion between exponents of the different approaches to the explanation of behavior, where the Bayesian defends his position by these arguments (Selten 1989: 5, 11, 21). Selten introduces several counterarguments, which, however, seem to be based on the assumption that Bayesian rationality has content.

Brian and Priscilla has nothing to do with the strategies they pick and, consequently, nothing with their success. The difference is just the extent to which their choices can be rationalized in terms of their beliefs. In fact, both might choose the same actions under identical circumstances. It is not as if the rules of Bayesianism offered protection against *mistakes* that could be identified as such by clever Priscilla. To Priscilla, there are actions that are mistaken in the light of one's beliefs but no actions that are mistaken just in light of the known facts. Priscilla is able to rationalize any behavior; even if Brian were unable to do the same, she could do it for him.

Thus, while there might be selective pressure to avoid certain kinds of behavior, there cannot be any selective pressure in favor of perfecting the rationalization of behavior. Nobody stands a better chance in any competition just on account of being a Bayesian. Of course, among Bayesians, there might be selective pressure against certain priors. But this is a completely different point.

The idea that there is selective pressure in favor of perfect rationality is historically connected with the as-if defense of the rationality postulate. The as-if defense has been the war cry of empiricist positive economics: "Never mind how people actually think; when it comes to action, they behave *as if* they were rational." The anything-goes theorem robs the as-if defense of any empiricist appeal, at least if rationality is taken to be Bayesian rationality. The statement "people behave as if they were Bayesians" turns out to be analytic; it boils down to "people do whatever they do." The adaptationist argument has been used to defend the as-if argument: Why should people behave as if they were rational? The adaptationist answer: Because those who do are more successful in the long run than those who do not. Obviously, this argument, if applied to Bayesian rationality, wrongly presupposes that Bayesian rationality helps to avoid mistakes.

6.3 Logic, Coherence, and the Axioms

Almost everybody agrees that deductive logic and logical consistency are valuable. However, deductive logic restricts only the structure of beliefs and not the choice of strategies. One might argue, then, that it cannot be held against Bayesianism if the Bayesian logic of probabilistic beliefs and the corresponding notion of consistency (often called "coherence") display the same weakness.[26]

This argument, however, is insufficient to defend Bayesianism. Logical consistency serves a purpose. Beliefs cannot possibly be true if they are inconsistent. Thus, if one wants truth, logical consistency is necessary. An analogous argument

[26] Analogies between Bayesianism and deductive logic are stressed by Howson (1997).

in favor of Bayesianism would have to point out some advantage of coherence unavailable to those relying on non-probabilistic beliefs and deductive logic alone. Such an argument is missing.

It is true that Bayesianism provides a logic of beliefs that rational persons must respect—*if and only if* their beliefs take the form of subjective probabilities. Those who reject this view—Popperians, classical statisticians, and others—can always point out to a Bayesian that their procedures could be rationalized on Bayesian grounds. The argument that the rationalizing prior might be "bad"—e.g., assign a positive probability to hypotheses they do not consider in earnest—will not worry them since for them there are no good priors anyway. In their view, beliefs just do not take the form of a probability distribution.

Of course, it is a theorem that the beliefs of rational agents should take the form of probabilities, not an assumption. The theorem follows from the axioms for preferences on the set of all strategies. If one argues that beliefs need not take the form of probabilities, the real question from a Bayesian point of view is, What's wrong with the axioms?

In my opinion, the axioms are quite reasonable if one is looking for a complete preference order on the set of all strategies. In the case of decision making under certainty or risk, a complete preference order might indeed be helpful.[27] But this is different in the case of choice under uncertainty.

Consider the following thought experiment. Adam has to choose from a menu in a restaurant. He orders chicken because he prefers it most. However, chicken is out. If pork was the second-best choice before, it is the best choice under the new conditions. A complete preference order means that he has decided in advance what to order if items are deleted from the menu.

There is no corresponding thought experiment for choice under uncertainty. Bayesian coherence amounts to a preference order among all conceivable strategies, where each strategy specifies the reactions to any new information. Hence, the reaction to the information that something is not available after all is already part of the chosen strategy. The preference order among the discarded strategies is therefore irrelevant by definition of the term strategy. Bayesianism not only is no help in choosing a strategy; it additionally requires that one chooses an order among the remaining irrelevant strategies. From this point of view, Bayesianism is worse than useless.

[27]Even then, rational choice does not require a complete preference order. Look at a simple example. Which of the following options would you prefer: losing your left hand, your right foot, or $ 10? I find the choice easy although completely ordering the alternatives is beyond me.

Acknowledgements

For useful hints, comments, and critical remarks on previous versions, I am grateful to Friedrich Breyer, Donald Gillies, Franco Gori, Daniel Hausman, Frank Heinemann, Ronald A. Heiner, Colin Howson, Karl-Josef Koch, Alberto Loayza-Grbic, Jürgen Meckl, Alan Musgrave, Dieter Schmidtchen, Reinhard Selten, Thusnelda Tivig, Heinrich Ursprung, Hermann Vetter, and Elie Zahar.

Bibliography

[Albert, 1996] M. Albert. Bayesian learning and expectations formation: Anything goes, Discussion Papers of the Faculty of Economics and Statistics, No 284, University of Konstanz, 1996.

[Albert, 1999] M. Albert. Bayesian learning when chaos looms large, *Economics Letters* 65, 1-7, 1999.

[Arrow, 1990] K. Arrow. Economic theory and the hypothesis of rationality. In Eatwell *et al.,* pp. 25-37, 1990.

[Anscombe and Aumann, 1963] F. Anscombe and R. J. Aumann. A definition of subjective probability, *Annals of Mathematical Statistics* 34, 199-205, 1963.

[Aumann, 1987] R. J. Aumann. Correlated equilibrium as an expression of Bayesian rationality, *Econometrica* 55, 1-18, 1987.

[Begg, 1982] D. K. H. Begg. *The Rational Expectations Revolution in Macroeconomics. Theories and Evidence,* Oxford: Phillip Allan, 1982.

[Bernstein, 1996] P. L. Bernstein. *Against the Gods. The Remarkable Story of Risk,* New York: Wiley, 1996.

[Bicchieri, 1993] C. Bicchieri. *Rationality and Coordination,* Cambridge: Cambridge University Press, 1993.

[Bicchieri *et al.,* 1997] C. Bicchieri, R. Jeffrey and B. Skyrms, eds. *The Dynamics of Norms,* Cambridge: Cambridge University Press, 1997.

[Binmore, 1993] K. Binmore. De-Bayesing game theory. In K. Binmore, A. Kirman and P. Tani, eds., *Frontiers of Game Theory,* pp. 321-339. Cambridge/Mass.: MIT Press, 1993.

[Binmore, 1994] K. Binmore. *Game Theory and the Social Contract, Vol. I: Playing Fair,* Cambridge/Mass.: MIT Press, 1994.

[Blume and Easley, 1995] L. E. Blume and D. Easley. What has the rational learning literature taught us?. In Kirman and Salmon, pp. 12-39, 1995.

[Border, 1992] K. C. Border. Revealed preference, stochastic dominance and choice of lotteries, *Journal of Economic Theory* 56, 20-42, 1992.

[Bray, 1983] M. Bray. Convergence to rational expectations equilibria. In Frydman and Phelps, 123-132, 1983a.

[Brémaud, 1988] P. Brémaud. *An Introduction to Probabilistic Modelling,* New York: Springer, 1988.

[Davies, 1987] P. Davis. *Cosmic Blueprint,* London: Heinemann, 1987.

[Devaney, 1989] R. L. Devaney. *An Introduction to Chaotic Dynamical Systems,* 2nd ed., Redwood City/Cal.: Addison-Wesley, 1989.

[Earman, 1992] J. Earman. *Bayes or Bust?,* Cambridge/Mass.: MIT Press, 1992.

[Eatwell *et al.,* 1990] J. Eatwell, M. Milgate and P. Newman, eds. *The New Palgrave: Utility and Probability,* New York: Norton, 1990.

[Ford, 1983] J. Ford. How random is a coin toss?, *Physics Today,* April 1983.

[Friedman, 1953] M. Friedman. The methodology of positive economics. In: M. Friedman, *Essays in Positive Economics,* pp. 3-43, Chicago and London: Chicago University Press, 1953.

[Frydman and Phelps, 1983] R. Frydman and E. S. Phelps. Introduction. In Frydman and Phelps, pp. 1-30, 1983a.

[Frydman and Phelps, 1983a] R. Frydman and E. S. Phelps. *Individual Forecasting and Aggregate Outcomes,* Cambridge: Cambridge University Press, 1983.

[Gillies, 2001] D. A. Gillies. Bayesianism and the Fixity of the Theoretical Framework. In present volume, pp. 373-390, 2001.

[Goldman, 1999] A. I. Goldman. *Knowledge in a Social World,* Oxford: Oxford University Press, 1999.

[Goodman, 1955] N. Goodman. *Fact, Fiction, and Forecast,* Indianapolis: Bobbs-Merrill, 1955.

[Green and Osband, 1991] E. J. Green and K. Osband. A revealed preference theory for expected utility, *Review of Economic Studies* 58, 677-696, 1991.

[Hacking, 1990] I. Hacking. Probability. In Eatwell *et al.,* pp. 163-177, 1990.

[Hahn, 1996] F. Hahn. Rerum cognoscere causas, *Economics and Philosophy* 12, 183-195, 1996.

[Howson, 1997] C. Howson. Logic and probability, *British Journal for the Philosophy of Science* 48, 517-531, 1997.

[Howson and Urbach, 1989] C. Howson and P. Urbach. *Scientific Reasoning: The Bayesian Approach,* La Salle/Ill.: Open Court, 1989.

[Kiefer and Nyarko, 1995] N. M. Kiefer and Y. Nyarko. Savage-Bayesian models of economics. In Kirman and Salmon, pp. 40-62, 1995.

[Kim, 1992] T. Kim. The subjective expected utility hypothesis and revealed preference, *Economic Theory* 1, 251-263, 1992.

[Kirman and Salmon, 1995] A. Kirman and M. Salmon, eds. *Learning and Rationality in Economics,* Oxford: Blackwell, 1995.

[Leamer, 1978] E. E. Leamer. *Specification searches,* New York: Wiley, 1978.

[Ledyard, 1986] J. O. Ledyard. The scope of the hypothesis of Bayesian equilibrium, *Journal of Economic Theory* 39, 59-82, 1986.

[Lindley, 1972] D. V. Lindley. *Bayesian Statistics. A Review,* Philadelphia: SIAM, 1972.

[Miller, 1994] D. Miller. *Critical Rationalism. A Restatement and Defense,* Chicago and La Salle/Ill.: Open Court, 1994.

[Musgrave, 1989] A. Musgrave. Saving science from scepticism. In F. D'Agostino and I. C. Jarvie, eds. *Freedom and Rationality,* pp. 297-323, Dordrecht: Kluwer, 1989.

[Muth, 1961] J. F. Muth. Rational expectations and the theory of price movements, *Econometrica* 29, 315-335, 1961.

[Nyarko, 1997] Y. Nyarko. Savage-Bayesian agents play a repeated game. In Bicchieri, Jeffrey and Skyrms, pp. 175-197, 1997.

[Pearl, 2000] J. Pearl. *Causality. Models, Reasoning, and Inference,* Cambridge: Cambridge University Press, 2000.

[Pesaran, 1987] M. H. Pesaran. *The Limits to Rational Expectations,* Oxford: Blackwell, 1987.

[Pope, 1991] R. E. Pope. The delusion of certainty in Savage's sure-thing principle, *Journal of Economic Psychology* 12, 209-241, 1991.

[Samuelson, 1947] P. A. Samuelson. *Foundations of Economic Analysis,* enlarged edition, Cambridge/ Mass. and London: Harvard University Press, 1983 (first published 1947).

[Savage, 1954] L. J. Savage. *The Foundations of Statistics,* New York: Wiley, 1954.

[Savage, 1962] L. J. Savage. *The Foundations of Statistical Inference,* London: Methuen/ New York: Wiley, 1962.

[Schuster, 1988] H. G. Schuster. *Deterministic Chaos,* 2nd rev. ed., Weinheim: VCH, 1988.

[Selten, 1989] R. Selten. *Evolution, Learning, and Economic Behavior,* 1989 Nancy L. Schwartz Memorial Lecture, J. L. Kellogg Graduate School of Management, Northwestern University.

[Simon, 1976] H. A. Simon. From substantive to procedural rationality. In S. J. Latsis, ed. *Method and Appraisal in Economics,* pp. 129-148, Cambridge: Cambridge University Press, 1976.

[Varian, 1992] H. R. Varian. *Microeconomic Analysis,* 3rd rev. ed., New York: Norton, 1992.