On the Rationality of Some Decision Rules in a Stochastic Environment

GUY COHEN AND PIERRE BERNHARD

***Abstract*—A classical decision rule consists of finding the decision which minimizes the expected cost. Liberty and Hartwig [1] proposed another decision rule which consists of minimizing a combination of the expectation and the variance of the cost in order to reduce the probability of bad realizations. We show that this decision rule does not meet a "minimal rationality requirement" in general. We relate it to another one and exhibit special cases when it does so.**

I. Introduction

This paper is motivated by the following consideration which appeared in a recent paper by Liberty and Hartwig [1]. Assuming a probabilistic description of the uncertainties, a classical decision rule is to choose the strategy which minimizes the expected cost. However, Liberty and Hartwig noticed that this way of looking at the first statistical moment only does not *a priori* reduce the probability of getting a bad performance on *one* realization. Hence, they proposed to minimize a combination of the expectation and the variance of the performance index.

We show, by an elementary example, that this new decision rule does not meet a "minimal rationality requirement" that we first define. Then we show that this decision rule can be considered as an approximation of another decision rule which exhibits this property. This eventually leads us to look at some special cases when the decision rule proposed by Liberty and Hartwig meets the requirement.

Manuscript received May 17, 1979.

G. Cohen is with the Centre d'Automatique–Informatique, Ecole Nationale Superieure des Mines de Paris, Fontainebleau, France.

P. Bernhard is with the Université de Paris IX-Dauphine, Paris, France, and the Centre d'Automatique–Informatique, Ecole Nationale Supérieure des Mines de Paris, Fontainebleau, France.

II. The Minimal Rationality Requirement

We consider a probability space $(\Omega, \mathcal{G}, P)$ where Ω is interpreted as the set of "states of nature." Generally speaking, there are two kinds of situations: P is either an "*objective*" or a "*subjective*" probability law. The former case is when there is some general agreement or statistical experience on P. In the latter, the decision maker has to express his personal feeling through the choice of P in order to be able to use the Bayesian approach which consists in the minimization of the expected cost. Let $(\mathcal{U}, \mathcal{B})$ be the measurable space of *decisions* and $(\mathcal{Y}, \mathcal{C})$ be the measurable space of *observations*. The *observation process* is described by a measurable mapping y from $\mathcal{U} \times \Omega$ to $\mathcal{Y}$. The *strategies* are (possibly a subset of) the set $\mathbb{S}$ of measurable mappings from $\mathcal{Y}$ to $\mathcal{U}$. Once a strategy s is chosen, to each realization ω a decision u is associated which is the solution of the implicit equation

$$s(y(u,\omega)) = u.$$

On the existence and uniqueness of this solution, that we shall hereafter denote by $u_s(\omega)$, see the causality conditions of Witsenhausen [2].

Let J be a measurable ($\mathbb{R}$-valued) functional on $\mathcal{U} \times \Omega$ which represents the performance index. Defining a "*decision rule*" consists of defining an $\mathbb{R}$-valued index μ on the admissible strategies in $\mathbb{S}$. Then the best decision will be that which minimizes $\mu(s)$.

The classical "*expected cost*" decision rule is given by

$$\mu_e(s) \triangleq \mathbb{E}J_s = \int_\Omega J(u_s(\omega),\omega)P(d\omega). \tag{1}$$

Liberty and Hartwig [1] suggested the decision rule

$$\mu_\alpha(s) \triangleq \mathbb{E}J_s + \alpha \operatorname{var} J_s \tag{2}$$

where α is a positive number and var denotes the variance (we assume the second-order moment exists).

A decision rule defines a preference order on admissible strategies, namely,

$$s_1 \overset{\mu}{\geqslant} s_2 \quad \text{iff } \mu(s_1) \geqslant \mu(s_2) \tag{3}$$

(and the same for strict inequalities).

Definition: A decision rule μ meets the "*minimal rationality requirement*" iff

$$\left[\Pr[J_{s_1} \geqslant J_{s_2}] = 1\right] \Rightarrow s_1 \overset{\mu}{\geqslant} s_2. \tag{4}$$

In the case when P is only a subjective probability law but assuming that there is a general agreement that Ω is the set of all possible states of nature, we can define a weaker minimal rationality requirement, independent from P, namely,

$$\left[\forall\omega\in\Omega,\ J(u_{s_1}(\omega),\omega) \geqslant J(u_{s_2}(\omega),\omega)\right] \Rightarrow s_1 \overset{\mu}{\geqslant} s_2. \tag{5}$$

The latter is useful when no P is introduced, e.g., for the *worst case* decision rule

$$\mu_{\max}(s) \triangleq \max_\omega J(u_s(\omega),\omega). \tag{6}$$

Clearly, any reasonable decision rule must meet either version of the minimal rationality requirement. We now show that this is not the case, in general, for that introduced by Liberty and Hartwig [1] [see (2)]. We first notice that (2) has to be modified in order to produce a preference order independent from the unit in which J is measured. As a matter of fact, if J is changed into λJ ($\lambda > 0$), then $\mathbb{E}J$ is changed into $\lambda\mathbb{E}J$ while var$J$ is changed into λ^2var$J$. We thus replace (2) by

$$\mu_\alpha(s) \triangleq EJ_s + \alpha\sigma J_s \tag{7}$$

where $\sigma J = (\text{var}J)^{1/2}$.

Example: Consider $\Omega = \{\omega_1, \omega_2\}$, both ω_i's with probability 0.5, $\mathcal{U} = \{u_1, u_2\}$, and suppose that the admissible strategies are reduced to constant mappings (open-loop controls) so that they can be identified to the decisions u_1, u_2. The performance index J is given by the following table:

	u_1	u_2
ω_1	0	21
ω_2	20	21

For $\alpha > 1.1$, u_2 will be preferred to u_1 although this is completely illogical.

III. Another Decision Rule

We consider the decision rule

$$\mu_\epsilon(s) \triangleq \inf\{a;\ \Pr[J_s \leq a] \geq 1-\epsilon\} \tag{8}$$

where ϵ is a given positive number. This can be interpreted as the accepted risk that the cost exceeds $\mu_\epsilon(s)$. A particular case of interest is for $\epsilon = 0$, when $\mu_\epsilon(s)$ turns out to be

$$\mu_0(s) = \operatorname*{ess\,sup}_{\omega} J(u_s(\omega), \omega). \tag{9}$$

If Ω is a discrete set of points ω_i and if for all i, $P(\omega_i)$ is nonnull, or if the mapping $\omega \to J(u_s(\omega), \omega)$ is continuous over the topological space Ω and if $P(\mathcal{O})$ is nonnull for all open subset $\mathcal{O}$ of Ω, then (9) coincides with (6). Hence, the decision rule (8) appears as a "softened" version of the worst case approach.

Lemma: The decision rule (8) has the property (4) [and *a fortiori* (5)].

Proof: The left-hand side of (4) implies that $\Pr[J_{s_1} \leq a] \leq \Pr[J_{s_2} \leq a]$ for all a, which, from (8), yields

$$\forall a:\ \mu_\epsilon(s_1) \leq a \Rightarrow \mu_\epsilon(s_2) \leq a$$

which implies the right-hand side of (4). □

The connection between (7) and (8) stems from the Bienaymé–Chebyshev inequality (assuming, for the time being, that $\sigma J_s > 0$)

$$\Pr[|J_s - EJ_s| \geq \alpha\sigma J_s] \leq \alpha^{-2}.$$

Hence,

$$\Pr[J_s \leq \mu_\alpha(s)] \geq 1 - \alpha^{-2} \tag{10}$$

where $\mu_\alpha(s)$ is defined by (7). Hence, $\mu_\alpha(s)$ appears as an overestimated value of $\mu_\epsilon(s)$ defined by (8) with $\epsilon = \alpha^{-2}$. We now make the following assumption.

Assumption: There exists an $\mathbf{R}$-valued function Π defined on $\mathbf{R}$ such that, for all admissible strategy s such that $\sigma J_s > 0$,

$$\forall \alpha \in \mathbf{R}:\ \Pr[J_s \leq EJ_s + \alpha\sigma J_s] = 1 - \Pi(\alpha). \tag{11}$$

The nontrivial fact in this assumption is that Π is independent of the strategy s, provided that $\sigma J_s > 0$ (otherwise, (11) is met with $\Pi(\alpha) \equiv 0$). Examples when such an assumption is met are: 1) Ω has only two possible outcomes (as in the example above) and 2) J_s is a Gaussian random variable if $\sigma J_s > 0$.

Proposition: Let

$$\hat{\alpha} = \inf\{\alpha;\ \Pi(\alpha) \leq \epsilon\} \tag{12}$$

where $\epsilon < 1$. Then $\hat{\alpha}$ is finite. Moreover the decision rule μ_ϵ [see (8)] and $\mu_{\hat{\alpha}}$ (see (7) with $\alpha = \hat{\alpha}$) coincide, i.e., $\mu_\epsilon(s) = \mu_{\hat{\alpha}}(s)$ for all admissible s.

Proof: First notice that for those s such that $\sigma J_s = 0$, we have $\mu_\epsilon(s) = \mu_\alpha(s) = EJ_s$ for all α and $\epsilon < 1$. Considering from now on that $\sigma J_s > 0$ we see from (10) and (11) that $\Pi(\alpha) \leq \alpha^{-2}$; thus $\hat{\alpha} \leq \epsilon^{-1/2} < +\infty$. Moreover, since $\lim \Pi(\alpha) = 1$ when $\alpha \to -\infty$, $\hat{\alpha} > -\infty$ when $\epsilon < 1$. Setting $a = \mu_\alpha(s) = EJ_s + \alpha\sigma J_s$, definition (8) can be rewritten as

$$\mu_\epsilon(s) = \inf_\alpha \{\mu_\alpha(s);\ \Pi(\alpha) \leq \epsilon\} = \mu_{\hat{\alpha}}(s),$$

the latter equality stemming from (12) and the fact that $\mu_\alpha(s)$ is an increasing function of α. □

Corollary: Under the above assumption, decision rule (7) has property (4) if $\forall \beta < \alpha$, $\Pi(\beta) > \Pi(\alpha)$.

Remarks: 1) With $\alpha = 0$, μ_α coincides with μ_e [see (1)] and thus, still has the property.

2) If Π is strictly decreasing, the property holds for any α.

3) In the example of the previous section, Π exists and is given by $\Pi(\alpha) = \{1 \text{ if } \alpha < -1;\ 0.5 \text{ if } -1 \leq \alpha < 1;\ 0 \text{ if } 1 \leq \alpha\}$. Hence the sufficient condition does not hold for $\alpha > 1$.

References

[1] S. R. Liberty and R. C. Hartwig, "Design-performance-measure statistics for stochastic linear control systems," *IEEE Trans. Automat. Contr.*, vol. AC-23, pp. 1085–1090, Dec. 1978.

[2] H. S. Witsenhausen, "On information structure, feedback and causality," *SIAM J. Contr.*, vol. 9, pp. 149–160, May 1971.