

Bayesian Paradigm and Fully Probabilistic Design ^{*}

M. Kárný ^{*}

^{*} Department of Adaptive Systems, Institute of Information Theory
 and Automation, Academy of Sciences of the Czech Republic, P.O.
 Box 18, 182 08 Prague 8, Czech Republic, (e-mail:
 school@utia.cas.cz, http://as.utia.cz)

Abstract: This text provides background of fully probabilistic design of decision-making strategies and finds its position with respect to the standard Bayesian decision making.

Keywords: Bayesian decision making, fully probabilistic design

1. INTRODUCTION

There is a wide range of axiomatic formulations of decision making (DM) under uncertainty and incomplete knowledge, e.g. Savage [1954]. It seems, however, that none of them fits satisfactorily to closed decision loops in which the selected actions influence distributions describing them, cf. Fishburn [1970], part three. This text is an engineering attempt to fill the gap. The adjective “engineering” means that the overall picture is preferred over subtleties like measurability of various mappings: technical assumptions of this type are implicitly

assumed to be valid. The text serves primarily as a formalized justification of the fully probabilistic design (FPD) of decision-making strategies, Kárný [1996], Guy and Kárný [2005], Kárný and Guy [2006]. Also, the relationship of the FPD to the standard Bayesian DM is established.

A considered DM unit, called here participant, selects a sequence of actions $a^{1:t} \equiv (a_1, \dots, a_t)$, $a_t \in a_t^* \neq \emptyset$, with the aim to influence its environment, a thought of part of the real world. In connection with the faced DM, the participant considers observations $\Delta^{1:t}$, $\Delta_t \in \Delta_t^* \neq \emptyset$, of the environment together with others unobserved variables $x^{1:t}$, $x_t \in x_t^*$. The collection of these variables

$$Q \equiv \left(\Delta^{1:t}, a^{1:t}, x^{1:t} \right) \equiv \left(d^{1:t}, x^{1:t} \right), \text{ horizon } t \text{ given, (1)}$$

forms behavior of the closed loop made of the participant and its environment. Typically, the behavior consists of a finite sequence of finite-dimensional vectors with real or discrete-valued entries. The inspected theory should help in selecting the optimal strategy among available DM strategies $R \equiv R^{1:t}$ formed by sequences of DM rules

$$R_t : d^{1:t-1} \rightarrow a_t^*. \quad (2)$$

The optimal strategy “pushes” the closed-loop behavior as much as possible towards a desired closed-loop behavior. It supposes that the participant has partial preferences among behaviors. The addressed theory of its construction is put together as follows.

- The partial preferential ordering \triangleleft on possible behaviors $Q \in Q^*$ is characterized and quantified by a loss function

^{*} This work was supported in part by grants GA ČR 102/08/0567, MŠMT 2C06001.

$Z : Q^* \rightarrow [-\infty, \infty] \equiv \mathfrak{R}$. The inevitable assumptions are only made. Consequently, the loss function $Z(\cdot)$ that orders possible behaviors *a posteriori* is non-unique, Section 2.

- For a fixed strategy R , possible behaviors $Q \in Q^*$ are expressed as images of external unobserved influences, called uncertainties, $N \in N^* \neq \emptyset$. Uncertainties include anything what a priori prevents unambiguous determination the closed-loop behavior Q . The mapping

$$W_R : N^* \rightarrow Q^* \quad (3)$$

induces *a priori* ordering \triangleleft_{R^*} of strategies $R \in R^*$ as the image of the *a posteriori* ordering \triangleleft of behaviors $Q \in Q^*$. Due to the presence of uncertainties $N \in N^*$, the induced ordering \triangleleft_{R^*} of strategies $R \in R^*$ is partial even if the *a posteriori* ordering \triangleleft of behaviors $Q \in Q^*$ is complete.

For selecting the optimal DM strategy a complete ordering has to be defined on R^* . The optimal strategy is then defined as the most preferred strategy in terms of this ordering. In order to respect participant’s preferences, it has to be an extension of the ordering induced by the ordering of behaviors. The extension is made via a local functional T quantifying partial ordering of losses generated by a variety of *a posteriori* orderings, their representations and by possible strategies. The integral representation of T , Rao [1987], is given by a kernel Φ and a finite regular Borrel measure μ . The measure is recognized as a universal model of uncertainties $N \in N^*$ common to all DM tasks sharing them. Even then, the extension is not unique but good kernels Φ must avoid bad strategies, Section 3.

- Optional mappings Z, Φ appearing during quantitative characterization of the *a priori* ordering of strategies \triangleleft_{R^*} are restricted by conditions that make the FPD the proper alternative in selecting the best strategy, Section 4.
- The standard Bayesian DM is shown to be a strict subset of the FPD, Section 5.
- General properties of the advocated FPD are summarized in Section 6.

2. A POSTERIORI ORDERING OF BEHAVIORS

The participant is supposed to have a strict preferential ordering \triangleleft among behaviors $Q \in Q^*$. It is the binary relation \triangleleft on ordered pairs $({}^aQ, {}^bQ) \in Q^* \times Q^*$

$${}^aQ \triangleleft {}^bQ \text{ reads } {}^aQ \text{ is preferred against } {}^bQ. \quad (4)$$

Preferential interpretation implies that the ordering \triangleleft is to be *asymmetric* ${}^aQ \triangleleft {}^bQ \Rightarrow \neg({}^bQ \triangleleft {}^aQ)$; if aQ is preferred against bQ then bQ is not preferred to aQ . Desirable consistency of preferences restricts the ordering to be *transitive* $({}^aQ \triangleleft {}^bQ \wedge {}^bQ \triangleleft {}^cQ) \Rightarrow {}^aQ \triangleleft {}^cQ$. The ordering \triangleleft is generally *partial* one as the participants are often unable or unwilling to compare all pairs of possible behaviors. This is the key but realistic obstacle of the preference modelling. The incomparable pairs can be perceived as indistinguishable

$${}^aQ \sim {}^bQ \Leftrightarrow (\neg({}^aQ \triangleleft {}^bQ) \wedge \neg({}^bQ \triangleleft {}^aQ)). \quad (5)$$

As a rule, the relation \sim is, however, intransitive so that it cannot be taken as equivalence. The intransitivity is easily demonstrated when considering the ordering of two-dimensional integer-valued vectors $Q = [Q_1, Q_2]$ with ${}^aQ \triangleleft {}^bQ \Leftrightarrow {}^aQ_i < {}^bQ_i, i = 1, 2$. It is sufficient to think about ${}^aQ = [0, 0], {}^bQ = [1, 1]$ and ${}^cQ = [2, -1]$ as

$${}^aQ \sim {}^cQ \wedge {}^bQ \sim {}^cQ \wedge {}^aQ \triangleleft {}^bQ. \quad (6)$$

It is possible to use (5) and define the transitive preferential equivalence \approx on which (6) is "forbidden". The equivalence relation \approx introduces on Q^* equivalence classes. Their collection we denote Q_{\approx}^* . The ordering \triangleleft induces the preferential ordering \triangleleft_{\approx} on Q_{\approx}^* . Let $A, B \in Q_{\approx}^*$, then

$$A \triangleleft_{\approx} B \Leftrightarrow \exists ({}^aQ \in A \wedge {}^bQ \in B) \text{ such that } {}^aQ \triangleleft {}^bQ.$$

This ordering is partial as still may exist ${}^aQ \sim {}^cQ \wedge \neg({}^aQ \approx {}^cQ)$. The above example demonstrates this.

The set $\tilde{Q}_{\approx}^* \subset Q_{\approx}^*$ is said \triangleleft_{\approx} -dense if for any $A \triangleleft_{\approx} B, A, B \in Q_{\approx}^* \setminus \tilde{Q}_{\approx}^*$ there is $C \in \tilde{Q}_{\approx}^*$ such that $A \triangleleft_{\approx} C$ and $C \triangleleft_{\approx} B$. This notion is needed in the following basic proposition about numerical representation of ordering \triangleleft .

Proposition 2.1. (Existence of real-valued loss function). Let exist a countable \triangleleft_{\approx} -dense set in Q_{\approx}^* . Then, there is a real-valued function Z defined on Q_{\approx}^* such that

$$\begin{aligned} {}^aQ \triangleleft {}^bQ &\Rightarrow Z({}^aQ) < Z({}^bQ) \text{ and} \\ {}^aQ \approx {}^bQ &\Rightarrow Z({}^aQ) = Z({}^bQ). \end{aligned} \quad (7)$$

Proof: See Fishburn [1970], Proposition 3.2. \square

The loss function $Z(\cdot)$ described in the above proposition is by no means unique. The freedom can be restricted in a meaningful way by requiring it to be continuous with respect to topology generated by the ordering \triangleleft_{\approx} . Still its uniqueness can be obtained at too high price of unnecessary additional assumptions.

3. A PRIORI ORDERING OF STRATEGIES

For a fixed environment, let us consider any version of the loss function meeting (7) and express the behavior Q as the image of the considered strategy $R \in R^*$ and of the uncertainties $N \in N^* \neq \emptyset$, cf. (3). Then, the loss function can be expressed as a compound mapping of uncertainties on the (extended) real line \Re

$$Z_R : N^* \rightarrow \Re \text{ with } Z_R(N) \equiv Z(W_R(N)). \quad (8)$$

Considering all possible strategies $R \in R^*$, we get the set Z_{R^*} of real-valued functions of uncertainty $N \in N^*$

$$Z_{R^*} \equiv \{Z : N^* \rightarrow \Re, \exists R \in R^* \text{ such that } Z = Z_R \text{ cf. (8)}\}. \quad (9)$$

The value-wise defined ordering of functions in Z_{R^*} induces a "natural" partial ordering \triangleleft_{R^*} of strategies ${}^aR, {}^bR \in R^*$

$${}^aR \triangleleft_{R^*} {}^bR \Leftrightarrow Z_{\Re}({}^aR) = Z(W_{{}^aR}(N)) \leq Z(W_{{}^bR}(N)) = Z_{{}^bR}(N) \quad (10)$$

$\forall N \in N^*$, while the inequality (10) is sharp on a "non-negligible" subset of N^* . The term "non-negligible" is made more exact below. The strategy bR in (10) is called dominated by the strategy aR and any reasonable ordering introduced on R^* must not take it as the optimal one: its consequences are worse than those of aR irrespectively of inaccessible uncertainties. This is the key requirement on the constructed ordering of strategies.

Requirement 3.1. (Inadmissibility of dominated strategies). Let the loss function and environment be fixed and let \triangleleft_{R^*} be defined by (10). Then, any admissible completion of this ordering must not prefer a dominated strategy on any subset of ${}^aR^* \subset R^*$.

Remark 1. (The quest for objectivity). By demanding non-dominance on subsets ${}^aR^*$ of R^* , we try to make the ordering weakly dependent on the specific-case-dependent choice of the set of considered strategies. This reflects the wish to make the proposed methodology as objective as possible. The resulting methodology should serve to a whole range of a posteriori ordering on behaviors and should be weakly dependent on the chosen numerical representations. The same line is followed throughout.

Assuming existence of countable $\triangleleft_{R^* \approx}$ -dense set in R^* , Proposition 2.1 guarantees existence of a functional $T : Z_{R^*} \rightarrow \Re$ such that ${}^aR \triangleleft_{R^*} {}^bR \Rightarrow T(Z_{{}^aR}) < T(Z_{{}^bR})$ and ${}^aR \approx_{R^*} {}^bR \Rightarrow T(Z_{{}^aR}) = T(Z_{{}^bR})$.

For a systematic choice of the best strategy, we have to define complete ordering on R^* . In order to respect the given a posteriori ordering of behaviors, it has to be an extension of the discussed partial ordering \triangleleft_{R^*} . The key property of such an extension (denoted further on also \triangleleft_{R^*}) is formulated in the following proposition.

Proposition 3.1. (Admissible ordering). Let countable $\triangleleft_{R^* \approx}$ -dense set in R^* exist. Then, Requirement 3.1 is met iff \triangleleft_{R^*} is quantified by the functional T fulfilling

$$\begin{aligned} {}^aR \triangleleft_{R^*} {}^bR &\Leftrightarrow T(Z_{{}^aR}) < T(Z_{{}^bR}) \text{ and } {}^aR \approx_{R^*} {}^bR \\ &\Leftrightarrow T(Z_{{}^aR}) = T(Z_{{}^bR}). \end{aligned} \quad (11)$$

Proof: The implication \Leftarrow in (11) has to be only proved. For it, it is sufficient to observe that the optimum strategy is dominated if the optimization is restricted to a subset ${}^aR^* \subset R^*$ consisting of strategies violating isotonicity. \square

This makes us to *define* the complete ordering of strategies by (11) for a loss function Z and a functional T .

4. BASIS OF THE FPD

In order to get operational tool for the choice of the best strategy, we represent the functional T ordering completely strategies by (11) by exploiting integral representation of local functionals. It is reasonable to require the functional T to be continuous on the space of "nice" loss functions Z , generated by possible choices of Z -versions and by the considered DM strategies $R \in R^*$. In accordance with Remark 1, we want to make the result weakly dependent

on a specific loss function $Z(\cdot)$ chosen, i.e., on the specific choice of a posteriori preferences \triangleleft . Thus, we consider the functional T acting on functions of the same uncertainty $N \in N^*$, which arise by both varying possible strategies $R \in R^*$ and loss functions $Z \in Z^*$ determining different a posteriori orderings (4) of possible behaviors $Q \in Q^*$. It means that the functional T (11) acts on

$$Z_{R^*}^* \equiv \cup_{Z \in Z^*} \text{loss functions on } Q^* Z_{R^*}, \text{ see (8), (9). (12)}$$

The nice loss functions in $Z_{R^*}^*$ are assumed to belong to the space of continuous real-valued functions with a compact support in N^* , which is a locally compact Hausdorff space. The uniform norm $\|\cdot\|$ is well defined on these functions. The following representation is described, up to minor changes in notations, by Theorem 5, page 479, Rao [1987], where the highly technical proof can be found.

Proposition 4.1. (Representation of local functional). Let $T : Z_{R^*}^* \rightarrow \mathfrak{R}$ be a mapping such that:

- (1) (Sequential continuity) If $\{Z_n : n \geq 1\} \subset Z_{R^*}^*$ is a bounded point-wise convergent sequence, then $\{T(Z_n) : n \geq 1\} \subset \mathfrak{R}$ is Cauchy.
- (2) (Additivity)

$$T({}^aZ + {}^bZ) = T({}^aZ) + T({}^bZ) \text{ if } {}^aZ {}^bZ = 0. \quad (13)$$
- (3) (Bounded uniform continuity) For each $\varepsilon > 0$, $\gamma > 0$, there is a $\delta \equiv \delta_{\varepsilon, \gamma}$ such that if $\|{}^aZ\| < \gamma$, $\|{}^bZ\| < \gamma$, ${}^aZ, {}^bZ \in Z_{R^*}^*$ and $\|{}^aZ - {}^bZ\| < \delta$, then $|T({}^aZ) - T({}^bZ)| < \varepsilon$.

Then, the functional T is representable as the integral

$$T(Z) = \int_{N^*} \Phi(Z(N), N) \mu(dN), \text{ where (14)}$$

μ is a finite regular Borel measure on N^* and the kernel $\Phi : \mathfrak{R} \times N^* \rightarrow \mathfrak{R}$ satisfies the following conditions:

- (4) $\Phi(0, N) = 0$ and $\Phi(\cdot, N)$ is continuous for almost all (a.a.) $N \in N^*$,
- (5) $\Phi(x, \cdot)$ is μ -measurable for all $x \in \mathfrak{R}$, and,
- (6) for each $Z \in Z_{R^*}^*$, $\Phi(Z(N), N)$ is bounded for a.a. $N \in N^*$ and for any point-wise convergent sequence $\{Z_n : n \geq 1\} \subset Z_{R^*}^*$ $\{\Phi(Z_n, \cdot) : n \geq 1\}$ is Cauchy in the space of integrable functions with measure μ .

Conversely, if the pair (Φ, μ) satisfies the last three conditions and the functional T is defined by (14), then it meets the initial three conditions.

Note that the “non-negligible sets” quoted in connection with (10) are those of a positive measure μ .

The only interpretation-sensitive assumption of the used theorem is additivity (13) of T on functions with disjoint supports. It seems to be widely acceptable as

$$Z = \underbrace{\chi_A Z}_{{}^aZ} + \underbrace{(1 - \chi_A)Z}_{{}^bZ}, \quad A \subset N^*, \quad \chi_A \equiv \begin{cases} 1 & \text{on } A \\ 0 & \text{otherwise} \end{cases}$$

both interprets the meaning of the sum of loss functions and justifies the wish for this additivity.

Obviously, the functional of the form (14) is admissible in the sense of Proposition 3.1 on $Z_{R^*}^*$ iff the kernel Φ is increasing function of the first argument for a.a. second arguments. The complete ordering \triangleleft_{R^*} (11) remains the same if T is multiplied by any positive constant.

Thus, without a loss of generality, we can assume that $\mu(N^*) = 1$, i.e., that μ is a probabilistic measure. Using Fundamental theorem of probability, Corollary 7, page 155 in Rao [1987], we can express the functional (14) as the functional on behaviors determined by the considered decision strategy $R \in R^*$

$$T(Z_R) = \int_{Q^*} \Phi(Z(Q), W_R^{-1}(Q)) dF_R(Q), \text{ where (15)}$$

$F_R(Q)$ is the probability distribution function of Q and W_R^{-1} is the inverse image of W_R (3).

The representation of the a priori ordering \triangleleft_{R^*} of the DM strategies $R \in R^*$ by the functional $T(Z_R)$ has the following important methodological consequences:

- The representation of the functional T separates description of uncertainty, $\mu(dN)$ in (14), and its influence on a posteriori ordering of behaviors $Q \in Q^*$, expressed by values of the specific loss function $Z(Q)$. Consequently, $\mu(dN)$ can be safely interpreted as the objective description of uncertainty entering the closed loop formed by the considered environment and DM strategy.
- The measure μ describing the uncertainties $N \in N^*$ converts into the probability distribution function $F_R(Q)$ of the closed-loop behaviors $Q \in Q^*$. This distribution depends explicitly on the chosen DM strategy R .

Taking into account the item ••, all definitions of uncertainties $N \in N^*$ and of the corresponding mappings W_R giving the same $F_R(Q)$ are equivalent. Thus, the probability distribution function $F_R(Q)$ is the universal model of the closed loop.

• The strategies-ordering functional $T(Z_R)$ is the expected value of $\Phi(Z(Q), W_R^{-1}(Q))$. Generally, it depends both on the values of the loss function Z (ordering behaviors a posteriori) and the mapping W_R (3) projecting the uncertainties $N \in N^*$ on behaviors $Q = W_R(N)$. Taking into account equivalence of all pairs $(N, W_R(N))$ with the same pdf $F_R(Q)$, see item •••, we can assume that $W_R^{-1}(Q) = \mathcal{W}(F_R(\cdot), Q)$ for a function \mathcal{W} . Thus, the kernel in (14) can be given the form

$$\begin{aligned} \Phi(Z(Q), W_R^{-1}(Q)) &= \Phi(Z(Q), \mathcal{W}(F_R(\cdot), Q)) \\ &\equiv \Omega(Z(Q), F_R(\cdot), Q) \end{aligned} \quad (16)$$

with the function Ω increasing in the first argument and $\Omega(0, F_R(\cdot), Q) = 0$ for a.a. Q .

• The kernel Ω models interaction between uncertainty, projected into Q^* , which is described by $F_R(\cdot)$, and a posteriori observable loss $Z(Q)$. It models attitude of the participant to risk (neutral, risk prone, risk aware) or more generally, a non-trivial interactions between a posteriori consequences and their distribution. There are strong indications, Starmer [2000], that such a possibility is badly needed at least in risk-facing DMs.

• In order to avoid technicalities, we assume further on that the probability distribution function $F_R(Q)$ has the probability density function (pdf) $f_R(Q)$ with respect to a dominating measure denoted here dQ .

Generally, the kernel Ω depends on the whole probability distribution function $F_R(\cdot)$. It makes, however, sense to make the acquired loss and thus the ordering of strategies

dependent only on realized behaviors, i.e., to adopt a sort of likelihood principle (the localness in Bernardo [1979]).

Requirement 4.1. (“Likelihood” principle). The pdf $f_R(\cdot)$ enters the kernel Ω (16) value-wise, i.e., $\Phi(Z(Q), W_R^{-1}(Q)) = \Phi(Z(Q), \mathcal{W}(f_R(\cdot), Q)) \equiv \Omega(Z(Q), f_R(Q), Q)$. (17)

The above considerations define the optimal strategy as minimizer of the expected value

$$E[\Omega(Z(Q), f_R(Q), Q)] \equiv \int \Omega(Z(Q), f_R(Q), Q) f_R(Q) dQ.$$

The optimized strategy influences just the pdf $f_R(Q)$, which enters both the function Ω and – linearly – the expectation operator E . Let us stress that the occurrence of $f_R(Q)$ in Ω is non-standard and represents the key generalization brought by the treated problem formulation.

For a given $\Omega(\cdot)$ and $Z(\cdot)$, let us define the *ideal pdf* ${}^I f(Q)$ ${}^I f(Q) \equiv f_{OR}(Q)$ with ${}^OR \in \text{Arg} \inf_{R \in R^*} E[\Omega(Z(Q), f_R(Q), Q)]$. (18)

The presented results indicate that neither the kernel Ω nor the loss Z are unique. At the same time, all designs leading to the same closed-loop description, the same pdf $f_R(Q)$, are equivalent. The fully probabilistic design takes the last statement seriously and formulates the design as the selection of the DM strategy that makes the closed-loop pdf $f_R(Q)$ as close as possible to the ideal pdf ${}^I f(Q)$. Thus, the ordering of the strategies is not determined via a loss function $Z(Q)$ and the kernel $\Omega(Z(Q), f_R(Q), Q)$ but via the ideal pdf ${}^I f(Q)$ and some kernel $\Upsilon(f_R(Q), Q) \equiv \Omega(Z(Q), f_R(Q), Q)$. The optimal strategy minimizes

$$E[\Upsilon] \equiv \int \Upsilon(f_R(Q), Q) f_R(Q) dQ. \quad (19)$$

While the ideal pdf is specific to a specific decision problem, the kernel Υ is a technical tool. We would like to find class of kernels suitable to a wide class of DM problems. The following requirements seem to be reasonable and were inspired by a related problem in Bernardo [1979].

Requirement 4.2. (Universal kernel Υ).

- The kernel Υ is to guarantee that the given ${}^I f$ is the only unrestricted minimizer of (19).
- $\Upsilon(f_R(Q), Q) = 0$ for a.a. $Q \in Q^*$ for which $f_R(Q) = {}^I f(Q)$, i.e., the contribution of such behaviors to the optimized functional $E[\Upsilon]$ is zero.
- Re-scaling of the behaviors does not change $E[\Upsilon]$.

The following proposition shows that under this requirement there is a little freedom in choosing the kernel Υ .

Proposition 4.2. (Form of the Υ). Under Requirement 4.1 and for an arbitrary ideal pdf ${}^I f$, let the kernel Υ meet Requirement 4.2 and have continuous derivatives with respect to the first argument for a.a. Q . Then,

$$\Upsilon(f_R(Q), Q) = A \ln \left(\frac{f_R(Q)}{{}^I f(Q)} \right), \quad A > 0, \text{ i.e.,} \quad (20)$$

$$E[\Upsilon] = A \int f_R(Q) \ln \left(\frac{f_R(Q)}{{}^I f(Q)} \right) dQ = A \times$$

Kullback-Leibler divergence (KLD) of $f_R(Q)$ on ${}^I f(Q)$.

Proof: By taking variations of the minimized functional (19) over pdfs on Q^* , we get the necessary condition for minimum

$$f(Q) \frac{\partial}{\partial f_R(Q)} \Upsilon(f_R(Q), Q) + \Upsilon(f_R(Q), Q) = A > 0 \quad (21)$$

This condition has to be met for $f_R(Q) = {}^I f(Q)$ for which the second term on the left-hand side of (21) disappear.

For $A \neq 0$, it has the solution $\Upsilon(f_R(Q), Q) = A \ln(f_R(Q)) + B(Q)$. The requirement $\Upsilon({}^I f(Q), Q) = 0$ determines $B(Q) = -A \ln({}^I f(Q))$ uniquely.

For $A = 0$, $\Upsilon(f_R(Q), Q)$ have to be with arbitrary precision proportional to the quadratic form $(f_R(Q) - {}^I f(Q))^2$ for all $f_R(Q)$ sufficiently close to ${}^I f(Q)$. This quadratic form is not, however, scale invariant. \square

Remark 2. (On Proposition 4.2 and its conditions).

- The KLD, Kullback and Leibler [1951],

$$D(f || {}^I f) \equiv \int_{Q^*} f(Q) \ln \left(\frac{f(Q)}{{}^I f(Q)} \right) dQ \quad (22)$$

is an often used proximity measure of pdfs with a range of applications in DM, statistics and information theory. It has an exceptional position within a class of so called f-divergences studied by many authors, e.g., Vajda [1989].

- The minimization of the KLD of f on ${}^I f$ is the essence of the FPD. It was proposed in Kárný [1996] on heuristic basis and extended into a general form in Kárný and Guy [2006]. The current paper tries to make this basis more firm and to relate the FPD to standard Bayesian DM.

- The closed-loop description $f_R(Q)$ enters into the optimized functional in a non-linear way. This is a source of strength as well as weakness of the FPD discussed in subsequent sections. Related considerations of conditional expectation as, possibly non-linear, mapping can be found in Pfanzagl [1967].

- It is worth stressing that the multi-modal ideal pdf allows a straightforward quantification of multiple-aims, which otherwise is taken as a hard extension of the standard single-aim Bayesian paradigm.

- For behaviors having non-numerical parts, the scaling invariance, Requirement 4.2, can be replaced by requirement on invariance to sufficient-type mappings.

5. BAYESIAN DM IS A SPECIAL CASE OF THE FPD

First we verify that the FPD defines a suitable loss and kernel corresponding with Proposition 4.2, i.e., we show that the constructed optimal strategy is non-dominated with respect to this loss. Then, the proposition corresponding to the section title is presented.

The FPD minimizes the expectation of $Z_R(Q) \equiv \ln \left(\frac{f_R(Q)}{{}^I f(Q)} \right)$ over the space pdfs $\{f_R(Q), R \in R^*\}$. The key practical restriction on allowable ideal pdfs

$${}^I f(Q) = 0 \Rightarrow f_R(Q) = 0, \text{ for any } R \in R^*. \quad (23)$$

allows us to stay within the class of continuous loss functions. This restricts the set of strategies among which the optimal one can be found. Essentially, it puts hard constraints on the desirable behaviors. It may of course happen that the restricted set of strategies is empty.

If we define $N^* = Q^*$, the isotonic kernel $\Phi(Z_R(Q), Q) = Z_R(Q)$ and the measure $\mu(dQ) = f_R(Q) dQ$, we see that the last three items of Proposition 4.1 are met. Consequently, the KLD is the local isotonic functional on functions

$Z_R(\mathcal{Q}) = \ln \left(\frac{f_R(\mathcal{Q})}{I_f(\mathcal{Q})} \right)$, $R \in R^*$, and meets the initial three conditions of Proposition 4.1.

For demonstrating the main result of this section, we reduce (17) to a special case

$$\Phi(Z(\mathcal{Q}), W_R^{-1}(\mathcal{Q})) = \Omega(Z(\mathcal{Q}), \mathcal{Q}) \quad (24)$$

with Ω having properties of the kernel Φ . Under this assumption, the local functional (15) reduces to the ordinary expectation of the isotonic transformation (via Ω) of the loss $Z(\cdot)$. In other words, the design reduces to the standard Bayesian one, which is characterized by linearity with respect to the optimized closed-loop model $f_R(\mathcal{Q})$ and a sort of “neutrality” with respect to this pdf.

Thus, the crucial question is whether we can select the ideal pdf $I_f(\mathcal{Q})$ such that – for a given environment, a given “neutral” kernel $\Omega(Z(\mathcal{Q}), \mathcal{Q})$ and a given loss function $Z(\mathcal{Q})$ – the FPD reduces to the standard Bayesian DM. In other words, we ask whether the FPD extends the Bayesian DM. The answer is affirmative.

Taking into account structure of the behavior (1), we can factorize the pdf $f_R(\mathcal{Q}) \equiv f_R(d^{1:\hat{t}}, x^{1:\hat{t}})$ describing it as

$$f_R(\mathcal{Q}) = \prod_{t=1}^{\hat{t}} \underbrace{f(\Delta_t, x_t | a^{1:t}, \Delta^{1:t-1}, x^{1:t-1})}_{\text{model of environment and its observation}} \times \underbrace{f(a_t | d^{1:t-1}, x^{1:t-1})}_{\text{rule } R_t \text{ of the strategy } R} \quad (25)$$

with $d^{1:0}, x^{1:0}$ representing trivial conditioning. The first generic factor in (25) is fixed when the environment is fixed. The second generic factor describes the optimized strategy. By definition, $x^{1:\hat{t}}$ are never directly observed by the participant. Thus, natural conditions of DM, Peterka [1981], have to be met

$$f(a_t | d^{1:t-1}, x^{1:t-1}) = f(a_t | d^{1:t-1}) \quad (26)$$

that comply with domains of decision rules R_t in (2)

Within the considered class of continuous loss functions on a compact support, the function

$$\mathcal{Z}(\mathcal{Q}) \equiv \Omega(Z(\mathcal{Q}), \mathcal{Q}), \text{ see (24)}, \quad (27)$$

is bounded. Thus, we can assume without a loss of generality that $\mathcal{Z}(\mathcal{Q})$, whose expectation is minimized in the standard Bayesian way, is bounded and non-negative. Moreover, $\mathcal{Z}(\mathcal{Q})$ can be always written in the additive form

$$\mathcal{Z}(\mathcal{Q}) = \sum_{t=1}^{\hat{t}} \underbrace{z(d^{1:t}, x^{1:t})}_{\text{partial loss}}, \quad z(d^{1:t}, x^{1:t}) \begin{cases} \mathcal{Z}(\mathcal{Q}) & \text{if } t = \hat{t} \\ 0 & \text{if } t < \hat{t} \end{cases} \quad (28)$$

Dynamic programming methodology, e.g. Bertsekas [2001], splits the optimization into a sequence of minimizations of the conditional expectation of the partial loss increased by the non-negative Bellman function of the same argument. Consequently, it is sufficient to express a static DM task as the FPD in order to get the conclusion valid for the general dynamic case.

Proposition 5.1. (Inclusion of Bayesian DM into the FPD). Let the bounded loss function $\mathcal{Z}(y, a) \geq 0$, $y \equiv (\Delta, x)$, and the environment model $f(y|a)$ be given. Let us search for an optimal randomized decision rule $O_f(a)$

$$O_f(a) \in \text{Arg min}_{f(a)} E[\mathcal{Z}]. \quad (29)$$

Let us exploit so called leave-to-fate option, Kárný et al. [2005], and define the ideal pdf $I_f(y, a) = I_f(y|a) I_f(a) \equiv I_f(y|a) f(a)$, i.e., the inspected strategy $f(a)$ is taken as the factor in the ideal pdf. The environment-related factor of the ideal pdf is chosen

$$I_f(y|a) = f(y|a) \exp \left\{ - \left[B(a) \mathcal{Z}(y, a) - (B(a) - 1) \int \mathcal{Z}(y, a) f(y|a) dy \right] \right\}, \quad (30)$$

where the non-negative coefficient $B(a)$ depends on the environment model and the loss function. It is found as a solution of the equation

$$1 = \int f(y|a) \exp \left\{ - \left[B(a) \mathcal{Z}(y, a) - (B(a) - 1) \int \mathcal{Z}(y, a) f(y|a) dy \right] \right\} dy, \quad (31)$$

which always exists. Then, the FPD with this ideal pdf reduces to the standard Bayesian DM, i.e.,

$$D(f || I_f) = E[\mathcal{Z}], \quad \forall \text{ strategies } f(a). \quad (32)$$

Proof: For a constant loss function, the assertion is trivial. Thus, the non-constant loss is generic. By construction, the factor $I_f(y|a)$ of the chosen ideal pdf (30) depends on the environment model and the loss function, i.e., it is independent of the decision rule $f(a)$ to be chosen. The condition (31) guarantees that it is a well defined pdf. Thus, it is just necessary to show that this condition can be met. Simple details are omitted to spare the space. \square

6. CONCLUSIONS

6.1 Drawbacks of the FPD

General limitations of the FPD follow from troubles with the KL divergence on pdfs having both continuous and Dirac-delta type constituents and thus having “tendency” to violate (23). Moreover, the optimized functional is non-linear and as such it is more sensitive to variations of its arguments f and I_f . The fact that aims are quantified in a non-standard way brings additional troubles:

- Expression of the real aims by I_f is non-trivial and it may happen that the option made does not reflect them.
- Some simple tasks like search for an extreme of a function are expressed in the FPD only in a cumbersome and sensitive way.
- The usual complete separation of the a posteriori loss and description of uncertainty is broken. This argument is, however, valid only when the neutral risk attitude is (implicitly) assumed.
- It is recommended to choose the ideal pdf as desirable and realistic modification of the environment model Kárný [2006]. While we take it as an advantage, there are tasks where this methodology brings undesirable effects.

6.2 Advantages of the FPD

At least because of authorship, it is not surprising that we let advantages of the FDP overweight disadvantages.

- Dynamic programming shows that stochastic optimization can be made by performing repetitively operations

(conditional expectation, minimization), Bertsekas [2001]. The FPD has an explicit minimizer, Kárný and Guy [2006], so that the (almost) inevitable approximation task, Si et al. [2004], is substantially simplified.

The evaluation complexity of the optimal strategy can be simply seen on so called data-driven FPD when no internal quantities are present and $\mathcal{Q} \equiv d^{1:\dot{t}}$. The following proposition is proved in Kárný et al. [2005].

Proposition 6.1. (Solution of the data-driven FPD). The optimal strategy minimizing the KLD of $f(\mathcal{Q}) \equiv f(d^{1:\dot{t}}) = \prod_{t=1}^{\dot{t}} f(\Delta_t | a_t, d^{1:t-1}) f(a_t | d^{1:t-1})$ on the ideal pdf $I f(\mathcal{Q}) \equiv I f(d^{1:\dot{t}}) = \prod_{t=1}^{\dot{t}} I f(\Delta_t | a_t, d^{1:t-1}) I f(a_t | d^{1:t-1})$

$$\text{is } f(a_t | d^{1:t-1}) = I f(a_t | d^{1:t-1}) \frac{\exp[-\omega(a_t, d^{1:t-1})]}{\gamma(d^{1:t-1})},$$

$$\gamma(d^{1:t-1}) \equiv \int I f(a_t | d^{1:t-1}) \times \exp[-\omega(a_t, d^{1:t-1})] da_t, \text{ for } t < \dot{t},$$

$$\omega(a_t, d^{1:t-1}) \equiv \int f(\Delta_t | a_t, d^{1:t-1}) \times \ln \left(\frac{f(\Delta_t | a_t, d^{1:t-1})}{\gamma(d^{1:t-1}) I f(\Delta_t | a_t, d^{1:t-1})} \right) d\Delta_t.$$

It is solved for $t = \dot{t}, \dot{t} - 1, \dots, 1$ with $\gamma(d^{1:\dot{t}}) = 1$.

Notice that the restricted support of the ideal pdf on actions implies restricted support of the chosen strategy. This manifests the general property of the FPD: the optimum lies between strategies meeting (23) and thus the ideal pdfs quantifies both decision aims and constraints.

- Multi-modal ideal pdf expresses “naturally” multiple decision aims, Böhm et al. [2005]. There is no conceptual jump between single and multiple aim optimization.
- In the multiple-participant context, the well-developed art of combining pdfs, Nelsen [1999], Cowell et al. [1999], Kárný et al. [2006], can be extended to combination of aims expressed by ideal pdfs Kracík [2004]. The similar problem is much harder in the classical setting.
- Unlike in the standard Bayesian DM, the optimal strategies are randomized. It is much more realistic as any channel implementing the designed DM strategy has a finite capacity, cf. Sims [2002], i.e., it is unable to implement non-randomized strategy.
- The reasonable choice of the ideal pdf as a modification of the current closed loop description, Kárný [2006], allows to respect the available environment model. For instance, quadratic optimization is performed when regulating linear environment with Gaussian noise while linear programming should be used for the same environment with uniform noise. This intuitively plausible result follows directly from the FPD problem formulation.

REFERENCES

J. M. Bernardo. Expected information as expected utility. *The Annals of Statistics*, 7(3):686–690, 1979.
 D.P. Bertsekas. *Dynamic Programming and Optimal Control*. Athena Scientific, Nashua, US, 2001. 2nd edition.

J. Böhm, T. V. Guy, and M. Kárný. Multiobjective probabilistic mixture control. In P. Horáček, M. Šimandl, and P. Zítek, editors, *Preprints of the 16th World Congress of the International Federation of Automatic Control*, pages 1–6, Prague, July 2005. IFAC.
 R. G. Cowell, A. P. Dawid, S. L. Lauritzen, and D. J. Spiegelhalter. *Probabilistic Networks and Expert Systems*. Springer-Verlag, 1999.
 P.C. Fishburn. *Utility Theory for Decision Making*. J. Wiley, New York, London, Sydney, Toronto, 1970.
 T. V. Guy and M. Kárný. Stationary fully probabilistic control design. In J. Filipe, J. A. Cetto, and J. L. Ferrier, editors, *Proceedings of the Second International Conference on Informatics in Control, Automation and Robotics*, pages 109–112, Barcelona, September 2005. INSTICC.
 Kárný. Choice of the ideal distribution for regulation by fully probabilistic design. *Automatica*, 2006. submitted.
 M. Kárný. Towards fully probabilistic control design. *Automatica*, 32(12):1719–1722, 1996.
 M. Kárný and T. V. Guy. Fully probabilistic control design. *Systems & Control Letters*, 55(4):259–265, 2006.
 M. Kárný, J. Böhm, T. V. Guy, L. Jirsa, I. Nagy, P. Nedoma, and L. Tesař. *Optimized Bayesian Dynamic Advising: Theory and Algorithms*. Springer, London, 2005.
 M. Kárný, J. Andryšek, A. Bodini, T. V. Guy, J. Kracík, and F. Ruggeri. How to exploit external model of data for parameter estimation? *International Journal of Adaptive Control and Signal Processing*, 20(1):41–50, 2006.
 J. Kracík. On composition of probability density functions. In J. Andryšek, M. Kárný, and J. Kracík, editors, *Multiple Participant Decision Making*, pages 113–121, Adelaide, May 2004. Advanced Knowledge International.
 S. Kullback and R. Leibler. On information and sufficiency. *Annals of Mathematical Statistics*, 22:79–87, 1951.
 R.B. Nelsen. *An Introduction to Copulas*. Springer, New York, 1999.
 V. Peterka. Bayesian system identification. In P. Eykhoff, editor, *Trends and Progress in System Identification*, pages 239–304. Pergamon Press, Oxford, 1981.
 J. Pfanzagl. Characterizations of conditional expectations. *The Annals of Mathematical Statistics*, 38(2):415–421, 1967.
 M.M. Rao. *Measure Theory and Integration*. John Wiley, New York, 1987.
 L.J. Savage. *Foundations of Statistics*. Wiley, New York, 1954.
 J. Si, A.G. Barto, W.B. Powell, and D. Wunsch, editors. *Handbook of Learning and Approximate Dynamic Programming*. Danvers, May 2004. Wiley-IEEE Press. ISBN 0-471-66054-X.
 C.A. Sims. Implications of rational inattention. Technical report, Department of Econometrics, Princeton University, 2002. 28 pages.
 C. Starmer. Developments in non-expected utility theory: The hunt for a descriptive theory of choice under risk. *Journal of Economic Literature*, XXXVIII:332–382, 2000.
 I. Vajda. *Theory of statistical inference and information*. Kluwer Academic Publishers, Dordrecht, 1989.