ELSEVIER

Operations Research Letters 21 (1997) 243–248

operations research letters

Bayesian optimization and genericity[1]

Christian Ewerhart

University of Bonn, Wirtschaftspolitische Abteilung, Adenauerallee 24-26, 53113 Bonn, Germany

Received 1 October 1996; revised 1 September 1997

Abstract

We show that in almost all decisions under uncertainty Bayesian optimization necessarily yields a decision which is optimal with respect to a full-support belief. © 1997 Elsevier Science B.V. All rights reserved.

Keywords: Bayesian decision theory; Genericity; Dominance; Normal-form games

1. Choice under uncertainty

A topic of vivid discussion in economic theory has always been the definition of rational behavior under uncertainty. The standard model here is that of the Bayesian decision maker who chooses a best reply to some subjective belief over the set of states of the world. What can be questioned, though, is whether rational individuals should hold priors that give positive probability to every state. Decision makers with priors of this kind seem to exhibit a more cautious behavior in the sense that they do not simply ignore the existence of certain contingencies. In this paper, we determine a generic class of decision situations with the property that for every decision situation in this class, it is true that every choice of a Bayesian decision maker can be interpreted as a choice of a cautious individual. More precisely, it is shown that, for every decision situation in this class, a mixed strategy is a best reply only if it is a best reply to some completely mixed prior.

Assume that there are n possible states of the world, denoted by $\omega_1, \ldots, \omega_n$ and collected in a set Ω. As usual, by a prior $p = (p_1, \ldots, p_n)$ we mean a probability distribution on Ω, where p_j denotes the probability assigned to state ω_j for every j. For every finite set A, let $\Delta(A)$ denote the set of probability distributions on A. In accordance with this notation, the set of priors is $\Delta(\Omega)$. A prior is said to be *completely mixed* if it gives positive probability to every state of the world. Let S denote the (finite) set of feasible choices for the decision maker. The decision maker receives utility $u(s, \omega)$ if s is his choice and ω turns out to be the true state of the world. Every decision situation can be specified by a *payoff matrix*. The rows of this matrix are indexed by the elements of S and its columns by the elements of Ω. Its entries are the numbers $u(s, \omega)$. We assume that the payoff matrix satisfies a certain genericity condition, which we call (G). This condition is defined and discussed in the next section. At this point it suffices to remark that (G) holds within an open and dense subset of the set of all payoff matrices.

[1] This paper is essentially Chapter 4 of my Ph.D. thesis in the European Doctoral Program.

PII S0167-6377(97)00050-3

In particular, the set of payoff matrices that does not fulfill (G) is of measure zero. We will allow mixed strategies, i.e. elements of $\Delta(S)$. Note that the utility function $u(\cdot,\cdot)$ is well-defined also on $\Delta(S)\times\Delta(\Omega)$, i.e. for every combination of mixed strategy and prior. We say that a mixed strategy $\sigma\in\Delta(S)$ is a *best reply* to a prior $p\in\Delta(\Omega)$ if $u(\sigma,p)\geqslant u(\sigma',p)$ for all $\sigma'\in\Delta(S)$.

Theorem 1. *For any decision situation which satisfies condition* (G), *the following is true*: *Every mixed strategy, which is a best response to some prior, is also a best reply to some completely mixed prior.*

The theorem says that every Bayesian rational choice in a generic decision situation can be interpreted as the choice of a cautious decision maker. It is possible to reformulate the result. For that, we recall some definitions. A mixed strategy $\sigma\in\Delta(S)$ is called *weakly dominated* if there exists a $\sigma'\in\Delta(S)$ such that $u(\sigma',p)\geqslant u(\sigma,p)$ for all $p\in\Delta(\Omega)$ and $u(\sigma',p)>u(\sigma,p)$ for at least one $p\in\Delta(\Omega)$. A *strictly dominated* strategy is a mixed strategy $\sigma\in\Delta(S)$ such that there exists a $\sigma'\in\Delta(S)$ with $u(\sigma',p)>u(\sigma,p)$ for all $p\in\Delta(\Omega)$. Theorem 1 now reads as follows.

Corollary 1. *There is no difference between weakly and strictly dominated mixed strategies for any player whose decision situation satisfies condition* (G).

Note that the analogous statement for *pure* strategies, where mixed strategy dominance is replaced by pure strategy dominance, is trivial.

Proof of Corollary 1. By a standard hyperplane argument it can be shown that a mixed strategy is a best reply if and only if it is not strictly dominated. Similarly, a mixed strategy is a best reply to a completely mixed belief if and only if it is not weakly dominated. A proof of this result can be found e.g. in Appendix B of Pearce [2]. □

2. The genericity assumption

At a first glance, one could suspect that all exceptional cases may be due to payoff ties, as in Fig. 1 below.

	ω_1	ω_2
s_1	1	1
s_2	1	0

Fig. 1. Payoff ties.

	ω_1	ω_2	ω_3
s_1	2	2	2
s_2	3	1	1
s_3	1	3	4

Fig. 2. s_1 is weakly but not strictly dominated.

In this decision situation, the choice s_2 is a best reply, but not to a completely mixed prior. However, excluding payoff ties does not suffice to insure that every rational choice is actually cautiously rational in the above defined sense. To see this, consider Fig. 2.

Here, the decision maker is never indifferent between two choices if he knows the state of the world. Nevertheless, there is a best reply that is not a best reply to any completely mixed prior. Specifically, the choice s_1 is a best reply to the prior that puts probability $\frac{1}{2}$ on ω_1 and ω_2, respectively, and probability 0 to ω_3. However, s_1 is never a best reply if the prior is completely mixed. For s_2 yields a higher expected payoff as long as the prior assigns a probability greater $\frac{1}{2}$ to ω_1 and s_3 performs better otherwise. Another way to describe the feature of this example is it to say that in spite of the fact that the decision maker is never indifferent between two outcomes, s_1 is weakly but not strictly dominated. The example is simple in the sense that there are obvious ways how to perturb the game in order to get rid of the weak dominance relation upon strategy s_1. What is not immediately apparent in general games, however, is whether the perturbation of the payoffs which removes a weak dominance relation upon one strategy does not lead to weak dominance upon another strategy. And even if this new weak dominance relation is eliminated by a second perturbation, this may induce yet another strategy to become weakly dominated. And so on. So in spite of the fact that a single non-genericity may quickly be removed it is not completely obvious if the same can be achieved for all strategies at one time. As could be seen above, it does not suffice to exclude ties among the payoffs. The genericity condition which will be imposed in the proof has the following geometric interpretation. To every choice of the decision maker there is a corresponding payoff vector in $\mathbb{R}^n$, where n denotes the number of states of the world. The convex hull of the payoff vectors that correspond to feasible choices is a polyeder in $\mathbb{R}^n$. A sufficient genericity condition is that, for all k, any k-dimensional facet of this polyeder is in general position to every single $(n-k)$-dimensional hyperplane that can be generated by vectors of the standard basis of the $\mathbb{R}^n$ alone (by definition, two subspaces U, V of the $\mathbb{R}^n$ are in general position if they span the whole space, i.e. if any vector may be represented as the sum of an element in U and of an element in V). This means that the facet when translated to the origin and extended to a subspace, intersects the hyperplane only in the origin. The condition becomes necessary and sufficient if we restrict the condition to facets in the “weak Pareto frontier” of the polyeder.

We will use the notion of a *coordinate vector* which is understood to be an element of the standard basis of the $\mathbb{R}^n$. E.g., if $n=3$, the coordinate vectors are $(1,0,0),(0,1,0)$, and $(0,0,1)$. Recall that n denotes the number of states of the world. Fix some $k \in \{1,\dots,n-1\}$ and consider the following condition.

(G_k) Whenever we choose $k+1$ arbitrary (pairwise different) row vectors $u_0,\dots,u_k$ from the payoff matrix and $n-k$ arbitrary (pairwise different) coordinate vectors $e_1,\dots,e_{n-k}$ of the $\mathbb{R}^n$, then

$$\det(e_1,\dots,e_{n-k},u_1-u_0,u_2-u_0,\dots,u_k-u_0) \neq 0. \quad (1)$$

Note that e_j is an arbitrary coordinate vector, and not necessarily the jth coordinate vector. Similarly, u_i does not necessarily correspond to the ith choice (we use this notation in order to avoid lowercase indices). We will say that assumption (G) is satisfied if for all $k \in \{1,\dots,n-1\}$, condition (G_k) holds. Before the general interpretation of the condition is given, we take a closer look at the case $k=1$.

Lemma 1 (Payoff ties). *A payoff matrix satisfies (G_1) if and only if the decision maker is never indifferent between two choices in any state of the world.*

Proof. By definition, (G_1) is the requirement that for any two choices and corresponding payoff vectors u_0 and u_1, and for any $n-1$ coordinate vectors $e_1,\dots,e_{n-1}$, the determinant $\det(e_1,\dots,e_{n-1},u_1-u_0)$ does not vanish. But this simply means that u_1 and u_0 do not have the same entry at the position that is not covered by the coordinate vectors $e_1,\dots,e_{n-1}$. □

So (G_1) requires that there must not be a column vector of the payoff matrix with two equal entries. It is therefore obvious that (G_1) has to be satisfied in order to derive the statement of Theorem 1. In the proof, we will impose not only (G_1), but also $(G_2),(G_3),\dots,(G_{n-1})$, where n is the number of states of the world. So if there are, say, three states of the world, there will be need to assume (G_2) in addition to (G_1). Condition (G_2) is not satisfied in the payoff matrix in Fig. 2. To see this, choose

$$u_0=\begin{pmatrix}2\\2\\2\end{pmatrix},\quad u_1=\begin{pmatrix}3\\1\\1\end{pmatrix},\quad u_2=\begin{pmatrix}1\\3\\4\end{pmatrix},\quad e_1=\begin{pmatrix}0\\0\\1\end{pmatrix}. \quad (2)$$

As $\det(e_1,u_1-u_0,u_2-u_0)=0$, the example in Fig. 2 does not satisfy (G_2). This shows that we cannot do without (G_2) in the case of three states of the world. We give one more example. Fig. 3 shows a decision situation that satisfies (G_1) and (G_2), but does not satisfy (G_3). This can be seen as follows. By Lemma 2, (G_1) becomes obvious. Checking (G_2) is the straightforward calculation of 24 determinants. To see that

	ω_1	ω_2	ω_3	ω_4
s_1	1	0	-3	7
s_2	0	1	5	-8
s_3	-1	-4	7	4
s_4	-2	-1	3	1

Fig. 3. A game that satisfies (G_1), (G_2), but not (G_3).

(G_3) does not hold, take

$$u_0 = \begin{pmatrix} 1 \\ 0 \\ -3 \\ 7 \end{pmatrix}, \qquad u_1 = \begin{pmatrix} 0 \\ 1 \\ 5 \\ -8 \end{pmatrix},$$

$$u_2 = \begin{pmatrix} -1 \\ -4 \\ 7 \\ 4 \end{pmatrix}, \qquad u_3 = \begin{pmatrix} -2 \\ -1 \\ 3 \\ 1 \end{pmatrix},$$

$$e_1 = \begin{pmatrix} 1 \\ 0 \\ 0 \\ 0 \end{pmatrix}.$$

Then $\det(e_1, u_1 - u_0, u_2 - u_0, u_3 - u_0) = 0$.

Consider choice s_4. It is a best reply to the prior that puts zero probability on ω_1, probability $\frac{6}{11}$ on ω_2, probability $\frac{3}{11}$ on ω_3, and probability $\frac{2}{11}$ on ω_4. Yet, it is weakly dominated by $1/3s_1 + 1/3s_2 + 1/3s_3$, hence not a best reply to a completely mixed prior by Lemma 1. As may have been noted, all examples given in this paper have an equal number of states and choices, respectively. The results clearly hold also in absense of this restriction.

Acknowledgements

I would like to thank Dieter Balkenborg, Tilman Börgers, Leonardo Felli, and Avner Shaked for helpful comments and encouragement. A grant provided by the German Academic Exchange Service is gratefully acknowledged.

Appendix

We prove that condition (G) implies the statement of Theorem 1. To get an intuition for the proof, recall that every choice of the decision maker may be considered as a point in the payoff space $\mathbb{R}^n$. The jth component of this point is the payoff in the case that the true state of nature is ω_j. In this geometric picture, the set of payoff vectors associated to all possible mixed strategies is a convex polytope in $\mathbb{R}^n$. Now take an arbitrary subset of the set of pure strategies that corresponds to a k-dimensional facet of this polytope. Then assumption (G) implies that the facet is in general position to all $(n-k)$-dimensional hyperplanes that can be generated by coordinate vectors only. This will rule out the possibility that a mixed strategy represented by a point in the facet is a best reply but not a best reply to a completely mixed prior. To prove the theorem, define the mapping $v : \Delta(S) \to \mathbb{R}^n$ that maps a mixed strategy σ to the vector $u(\sigma, \cdot)$ in the payoff space. The expected utility $u(\sigma, p)$ of a strategy σ with respect to a prior $p \in \Delta(\Omega)$ may then be calculated as the scalar product of $v(\sigma)$ and p, i.e. $u(\sigma, p) = v(\sigma) \cdot p$. The set of mixed best replies to a prior p will be denoted by

$$\mathrm{BR}(p) = \{\sigma \mid v(\sigma) \cdot p = \max!\}. \tag{3}$$

After these preparations, let $\bar{\sigma} \in \mathrm{BR}(\bar{p})$ be a best reply to some subjective prior $\bar{p}$ over the set of states of the world. The pointwise image of the simplex $\mathrm{BR}(\bar{p})$ under the linear mapping v is a convex polytope in the payoff space $\mathbb{R}^n$ and will be denoted by $v(\mathrm{BR}(\bar{p}))$. It is a facet of the polytope

$$v(\Delta(S)) = \mathrm{conv}\{v(s) \mid s \in S\}, \tag{4}$$

which itself is the convex hull of the set $\{v(s) \mid s \in S\}$ of row vectors of the payoff matrix. Hence, there exist strategies $s_0, \ldots, s_k \in S$ and corresponding row vectors $u_j = v(s_j)$ of the payoff matrix such that $v(\mathrm{BR}(\bar{p})) = \mathrm{conv}(u_0, \ldots, u_k)$. We assume k to be minimal with respect to this property. We will now define a subspace Y which is coplanar to (and of the same dimension as) the facet under consideration. For this, translate the polytope $v(\mathrm{BR}(\bar{p}))$ to the origin by subtracting the vector u_0. Let

$$Y = \langle v(\mathrm{BR}(\bar{p})) - u_0 \rangle \tag{5}$$

be the linear subspace of the payoff space that is generated by the translated polytope. Obviously,

$$Y = \langle u_1 - u_0, \ldots, u_k - u_0 \rangle . \tag{6}$$

The minimality requirement on k implies that the dimension of the vector space Y is precisely k. The orthogonal space of Y, i.e. the set

$$Y^{\perp} = \{x \in \mathbb{R}^n \mid x \cdot y = 0 \text{ for all } y \in Y\}, \tag{7}$$

will be of fundamental importance for the proof. Elements of $Y^{\perp}$ should be seen as candidate "beliefs", with respect to which $\bar{\sigma}$ could be interpreted as a best response. They will be called stabilizing vectors. All strategies $s_0, \dots, s_k$ are best replies to $\bar{p}$, hence they are payoff equivalent with respect to $\bar{p}$. So

$$v(s_j) \cdot \bar{p} = v(s_0) \cdot \bar{p}, \tag{8}$$

hence $u_j \cdot \bar{p} = u_0 \cdot \bar{p}$ for all $j = 1, \dots, k$. Therefore, $\bar{p} \in Y^{\perp} \setminus \{0\}$. In particular, $0 \leqslant k \leqslant n-1$, i.e. the facet $v(\mathrm{BR}(\bar{p}))$ cannot have full dimension in $\mathbb{R}^n$. Now choose a vector $w \in Y^{\perp} \setminus \{0\}$ with the following properties.

1. $w_i = 0$ for all i with $\bar{p}_i = 0$ and
2. among all w satisfying condition 1, w has a maximal number of zero components.

Notice that, since $\bar{p} \in Y^{\perp} \setminus \{0\}$, it is always possible to find such a w. The vector w has the following interpretation. There exists a minimal subspace generated only by vectors of the standard basis that includes the prior $\bar{p}$. In this subspace, we look for a stabilizing vector w which lies in a subspace of the same kind, again with minimal dimension. This subspace can be defined as follows. Let

$$H = \{x \in \mathbb{R}^n \mid x_i = 0 \text{ if } w_i = 0, \text{ for all } i\}. \tag{9}$$

H is a with respect to dimension smallest subspace of the payoff space with the property that it is generated by vectors of the standard basis and that it contains a non-trivial vector that is orthogonal to the facet $v(\mathrm{BR}(\bar{p}))$. First we show that the genericity assumption (G) implies $\dim(H) > k$. Recall that k is both the dimension of the facet $v(\mathrm{BR}(\bar{p}))$ and the dimension of Y. Assume to the contrary that $\dim(H) \leqslant k$. Since $\dim(H)$ is equal to the number of nonzero entries of w, we can assume without loss of generality that

$$w \in \langle e_1, \dots, e_k \rangle. \tag{10}$$

Then w is orthogonal to $\langle e_{k+1}, \dots, e_n \rangle$. But w is orthogonal to Y as well, and therefore Y and $\langle e_{k+1}, \dots, e_n \rangle$ cannot generate the whole of the $\mathbb{R}^n$. Consequently, since $Y = \langle u_1 - u_0, \dots, u_k - u_0 \rangle$, the determinant

$$\det(u_1 - u_0, \dots, u_k - u_0, e_{k+1}, \dots, e_n) = 0 \tag{11}$$

must vanish. This, however, contradicts the genericity condition. Hence the first assertion is proved. Next we show that $\dim(H \cap Y^{\perp}) \leqslant 1$. Assume that $\dim(H \cap Y^{\perp}) > 1$. Take any nonzero component w_j of w. Then the equation $x_j = 0$ has a nonzero solution w' in the at least two-dimensional linear space $H \cap Y^{\perp}$. However, the vector w' possesses more zero entries than w in contradiction to the maximality condition 2 above. This proves the second statement. We can start to construct the completely mixed prior we are looking for. For this, consider the orthoprojection $\pi : \mathbb{R}^n \to H^{\perp}$ whose kernel is precisely H. The linear space H is at least of dimension $k+1$, as we saw above. Consequently, the space $H^{\perp}$ is at most of dimension $n-k-1$. On the other hand, by the above statement about the dimension of $H \cap Y^{\perp}$,

$$\dim(\pi(Y^{\perp})) = \dim(Y^{\perp}) - \dim(Y^{\perp} \cap H) \geqslant n-k-1, \tag{12}$$

implying that $\pi(Y^{\perp}) = H^{\perp}$. In particular, there exist a $z \in Y^{\perp}$ such that $\pi(z)$ has only positive entries. The vector z may have negative entries at those positions where w has non-zero entries. However, if $w_i \neq 0$ then $p_i > 0$ by choice of w. Therefore, for a sufficiently small $\varepsilon > 0$, the expression $\varepsilon z + \bar{p}$ has only positive entries, too, and normalized such that its entries sum up to one it will be denoted by p_{ε}. Thus

$$p_{\varepsilon} = \frac{\varepsilon z + \bar{p}}{|\varepsilon z + \bar{p}|}, \tag{13}$$

where $|p| = |p_1| + \cdots + |p_n|$. Since $Y^{\perp}$ is a linear space, necessarily $p_{\varepsilon} \in Y^{\perp} \cap \Delta(\Omega)$. Finally it will be shown that if ε is chosen sufficiently small then $\bar{\sigma}$ is a best reply to the completely mixed belief p_{ε}. For this, it suffices to prove that the set

$$O = \{p \in Y^{\perp} \cap \Delta(\Omega) \mid \mathrm{BR}(p) = \mathrm{BR}(\bar{p})\} \tag{14}$$

is relatively open in $Y^{\perp} \cap \Delta(\Omega)$. The best reply correspondence has a closed graph (see [1], p. 30). Hence, for any pure strategy s, the set $\{p \in \Delta(\Omega) \mid s \in \mathrm{BR}(p)\}$ is closed. Consequently,

$$V_s = \{p \in \Delta(\Omega) \mid s \notin \mathrm{BR}(p)\} \tag{15}$$

is an open set for any s. It follows that the set

$$\tilde{O} = \{p \in \Delta(\Omega) \mid \mathrm{BR}(p) \subseteq \mathrm{BR}(\bar{p})\}, \tag{16}$$

being the intersection of the finitely many sets V_s with $s \notin \mathrm{BR}(\bar{p})$, is open in $\Delta(\Omega)$, too. It remains to be shown that $\tilde{O} \cap Y^{\perp} = O$. The inclusion $\tilde{O} \cap Y^{\perp} \supseteq O$ follows immediately from the definitions. For the reverse inclusion let $p \in \tilde{O} \cap Y^{\perp}$. We must show that

$\mathrm{BR}(p) \supseteq \mathrm{BR}(\bar{p})$. Let $\sigma \in \mathrm{BR}(\bar{p})$. Then, choose some $\sigma' \in \mathrm{BR}(p)$. Since $p \in \tilde{O}$, we know that actually $\sigma' \in \mathrm{BR}(\bar{p})$. But, as $p \in Y^{\perp}$ and $v(\sigma) - v(\sigma') \in Y$, we get $v(\sigma) \cdot p = v(\sigma') \cdot p$. Therefore, $u(\sigma, p) = u(\sigma', p)$, and consequently σ is a best reply to p, too. This shows $\mathrm{BR}(p) = \mathrm{BR}(\bar{p})$. Thus, $\tilde{O} \cap Y^{\perp} = O$. This completes the proof of Theorem 1. □

References

[1] D. Fudenberg, J. Tirole, Game Theory, MIT Press, Cambridge, MA, 1991.

[2] D.G. Pearce, Rationalizable strategic behavior and the problem of perfection, Econometrica 52 (1984) 1029–1050.