Synthese (2010) 172:157–176
DOI 10.1007/s11229-009-9470-7

Coherent choice functions under uncertainty

Teddy Seidenfeld · Mark J. Schervish ·
Joseph B. Kadane

Received: 10 June 2008 / Accepted: 13 August 2008 / Published online: 7 March 2009

Abstract We discuss several features of *coherent choice functions*—where the admissible options in a decision problem are exactly those that maximize expected utility for some probability/utility pair in fixed set ***S*** of probability/utility pairs. In this paper we consider, primarily, normal form decision problems under uncertainty—where only the probability component of ***S*** is indeterminate and utility for two privileged outcomes is determinate. Coherent choice distinguishes between each pair of sets of probabilities regardless the "shape" or "connectedness" of the sets of probabilities. We axiomatize the theory of choice functions and show these axioms are necessary for coherence. The axioms are sufficient for coherence using a set of probability/almost-state-independent utility pairs. We give sufficient conditions when a choice function satisfying our axioms is represented by a set of probability/state-independent utility pairs with a common utility.

Keywords Choice functions · Coherence · Γ-*Maximin* · *Maximality* · Uncertainty · State-independent utility

1 Introduction

In this paper we continue our study of coherent choice functions, which we started in our (Kadane et al. 2004) "Rubinesque" theory of decision. Let ***O*** be a set of feasible

T. Seidenfeld (✉) · M. J. Schervish · J. B. Kadane
Department of Philosophy, Carnegie Mellon University, Baker Hall, 135, Pittsburgh, PA 15213, USA
e-mail: teddy@stat.cmu.edu

M. J. Schervish
e-mail: mark@stat.cmu.edu

J. B. Kadane
e-mail: kadane@stat.cmu.edu

options, which we also call an *option set*. A ***choice function*** $\boldsymbol{C}(\boldsymbol{O})$ identifies the (possibly empty) subset of $\boldsymbol{O}$ that are the $\boldsymbol{C}$-admissible options in the decision problem given by the option set $\boldsymbol{O}$. We say that $\boldsymbol{C}(\bullet)$ is ***coherent*** provided that there is a non-empty set $\boldsymbol{S}$ of probability/utility pairs $\boldsymbol{S} = \{(\boldsymbol{p}, \boldsymbol{u})\}$ such that the $\boldsymbol{C}$-admissible options are precisely those that are Bayes with respect to some probability/utility pair $(\boldsymbol{p},\boldsymbol{u})$ in $\boldsymbol{S}$. That is, a coherent choice function satisfies the condition that for each feasible option $o \in \boldsymbol{O}$, o is $\boldsymbol{C}$-admissible, $o \in \boldsymbol{C}(\mathbf{O})$, if and only if there is a pair $(\boldsymbol{p}, \boldsymbol{u}) \in \boldsymbol{S}$ such that o maximizes the $\boldsymbol{p}$-expected $\boldsymbol{u}$-utility over $\boldsymbol{O}$. For short, we will call these the *Bayes-admissible options* in $\boldsymbol{O}$ with respect to $\boldsymbol{S}$.

Here we consider decision problems under uncertainty, where utility for two privileged outcomes is determinate. That is, in this paper, the target representation for a coherent choice function is a set of probability/utility pairs with a common cardinal utility function for two outcomes, which serve as the **0** and **1** of each utility function, $\boldsymbol{u}$. For simplicity, we restrict attention to option sets $\boldsymbol{O}$ where $\boldsymbol{C}$-admissible options exist. In order to assure that, i.e., so that $\boldsymbol{C}(\boldsymbol{O})$ is not empty we require that the option set be closed (as we make precise in Sect. 2). By contrast, if $\boldsymbol{O}$ is not closed, then given a set $\boldsymbol{S}$ there may be no Bayes-admissible options in $\boldsymbol{O}$. For a familiar example, if utility is increasing in the quantity X, then in a decision-under-certainty problem—where the decision maker chooses from an infinite menu of sure outcomes, and where probability is irrelevant—with $\boldsymbol{O} = \{0 \le x < 1\}$ each option in $\boldsymbol{O}$ is Bayes-inadmissible.

The use of a coherent choice function coincides with Levi's (1980) principle of E-admissibility in cases where the set S is a cross-product of a convex set of probabilities and a convex set of utilities: $\boldsymbol{S} = \boldsymbol{P} \times \boldsymbol{U}$—with convex sets $\boldsymbol{P}$ and $\mathbf{U}$. Also, we find that Savage [(1954, pp. 123–124), where he argues that option b is "superfluous" for the decision pictured by his Fig. 1] endorses a coherent choice rule with $\boldsymbol{S}$ a convex set of probabilities and a common utility.

We adopt the framework of choice functions, rather than using a binary preference relation because coherent choice (as used here) does not reduce to pairwise comparisons. The following example, which we repeat from our ISIPTA-03 paper, illustrates this theme.

Example 1 Consider a binary decision problem involving two states of uncertainty, $\Omega = \{\omega_1, \omega_2\}$ with three feasible options $\boldsymbol{O} = \{f, g, h\}$, and where utility is determinate: $\boldsymbol{u}(f(\omega_1)) = \boldsymbol{u}(g(\omega_2)) = 0.0$, $\boldsymbol{u}(f(\omega_2)) = \boldsymbol{u}(g(\omega_1)) = 1.0$, and $\boldsymbol{u}(h(\omega_1) = \boldsymbol{u}(h(\omega_2)) = 0.4$. Let uncertainty over the states be indeterminate, with $\boldsymbol{P} = \{\boldsymbol{p} : 0.25 \le \boldsymbol{p}(\omega_2) \le .75\}$. Figure 1 shows the graph of expected utilities for each option. Thick lines depict the surface of expected-utility maximization.

We rehearse three decision rules for use in this problem.

Γ-*Maximin*. Maximize minimum expected utility over the feasible options. This rule is well studied in Gilboa and Schmeidler (1989). In brief, Γ-*Maximin* induces a preference ordering over options, but fails the von Neumann–Morgenstern *Independence* postulate. Under Γ-*Maximin* only $\{h\}$ is admissible from the set $\{f, g, h\}$. In Fig. 1, this is evident as the low point of each of the graphs of f and of g is below that of the constant act h.

Maximality (Sen/Walley). Admissible options are those that are undominated in expectations (over all $\boldsymbol{p} \in \boldsymbol{P}$) by any single alternative option. Under *Maximality* all

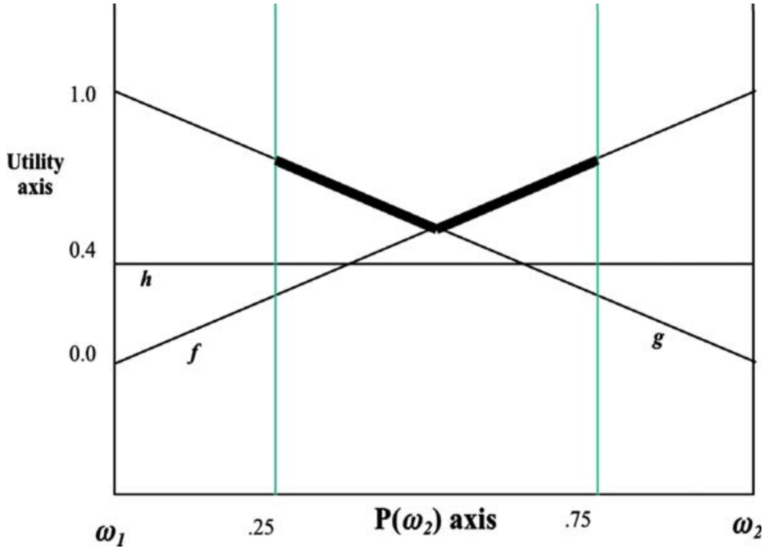

Fig. 1 The graph of expected utilities for each act in $\boldsymbol{O} = \{f, g, h\}$. The surface of Bayes-admissibility is bold

three options are admissible from the set $\{f, g, h\}$ as none dominates the others in $\boldsymbol{P}$-expectations under pairwise comparisons. In Fig. 1, this is seen by noting that the graphs for each pair of act cross each other. *Maximality* does not induce a preference ordering over options; nonetheless, admissibility is given by pairwise comparisons. Note that if the option set is expanded to include mixed options, then h is no longer maximal, since the "fair" mixture h' of f and g, denoted $.5f \oplus .5g$, which corresponds to a constant act with expected utility 0.5, has greater expected utility than does h for each $\boldsymbol{p} \in \boldsymbol{P}$. As is evident then, whether an option (e.g., option h) is admissible under *Maximality* depends upon whether the set of feasible options is closed under mixtures.

Coherent choice. Since the set of probabilities $\boldsymbol{P}$ is convex in this example, coherent choice reduces to Levi's rule of. E-admissibility—admissible choices have Bayes' models, i.e., they maximize expected utility for some probability in the (convex) set $\boldsymbol{P}$. Subset $\{f, g\}$ identifies the Bayes-admissible options from $\{f, g, h\}$ under *Coherent Choice*. In Fig. 1, the surface of Bayes-admissible options meets acts f and g, but not h. In this setting, Levi (1986, Sect. 5.2), calls the Bayes-inadmissible option h "second worst." E-admissibility does not induce an ordering over options and does not reduce to pairwise comparisons either, as the following illustrates.

Consider the three pairwise choice problems using binary subsets of $\boldsymbol{O} = \{f, g, h\}$. In such pairwise choices, both options are E-admissible, even though h is inadmissible in the feasible set of three options $\boldsymbol{O}$. But if h is replaced by a ("second best") constant option h' with utility $\boldsymbol{u}(h'(\omega_1)) = \boldsymbol{u}(h'(\omega_2)) = 0.5$. then with respect to the set $\boldsymbol{P}$, in the decision problem with feasible set $\boldsymbol{O}' = \{f, g, h'\}$, all three options are E-admissible and likewise, in pairwise choices between any two of them each is E-admissible. Hence, E-admissibility in pairwise choice does not determine E-admissibility from larger option sets. Pairwise choice is insufficient, generally, to distinguish second best from second worst.

As noted above, h is not "Bayes" in $\boldsymbol{O}$ with respect to $\boldsymbol{P}$. More dramatically, h is (uniformly) dominated by some mixtures of f and g. The constant mixed option $h' = .5f \oplus .5g$, with expected utility 0.5 independent of $\boldsymbol{p}$, uniformly dominates h. This is no coincidence, as the following result establishes.

Let $\Omega = \{\omega_1, \omega_2, \ldots, \omega_n\}$ be a finite partition of states. Let $\boldsymbol{O} = \{o_1, o_2, \ldots, o_m\}$ be a *finite* set of options defined on Ω, such that for $\boldsymbol{o_i} \in \boldsymbol{O}$, $\boldsymbol{u}(o_i(\omega_j)) = \boldsymbol{u_{ij}}$, a determinate cardinal utility of the consequence of o_i when state ω_j obtains. Let $\boldsymbol{P}$ be the class of all probability distributions over Ω. Similarly, let $\boldsymbol{Q}$ be the class of all (simple) mixed acts over $\boldsymbol{O}$, with a mixed act denoted q.

Theorem 1 (Pearce 1984, p. 1048) *Suppose for each $\boldsymbol{p} \in \boldsymbol{P}$, act $o^* \in \boldsymbol{O}$ fails to maximize expected utility. Then there is a mixed alternative q^* that (uniformly) strictly dominates o^*. That is, $\boldsymbol{u}(q^*(\omega_j)) > \boldsymbol{u}(o^*(\omega_j)) + \varepsilon$, for $j = 1, \ldots, n$, with $\varepsilon > 0$.*

With this result we are able to apply the strict standard of de Finetti's "incoherence" (= uniform, strict dominance), to a broad class of decisions under uncertainty, analogous to the scope of traditional *Complete Class* theorems for Bayes decisions (Wald 1950). That is, de Finetti (1974) uses his theory of coherent previsions in a class of decision problems where the decision maker's sole options are to fix *fair* gambling rates over random variables assuming a linear utility in payoffs. As explained in Sect. 2, the domain of our coherent choice functions are (simple) horse-lotteries. When the state-space is finite, this is a much larger class of decision problems than is addressed in de Finetti's theory. By contrast with Wald's theory, the standard of coherence used here—avoiding strict dominance—is more generous than Wald's requirement of *admissibility* (= avoiding weak dominance), as is used in his *Complete Class* theorems. That is, some Bayes-admissible options in a decision problem may be inadmissible in Wald's theory. But not so in the theory developed here.

2 Distinguishing sets of probabilities by their coherent choice functions

Consider a finite state space $\Omega = \{\omega_1, \ldots, \omega_n\}$ with the class of all options $\mathcal{H}$ the (Anscombe and Aumann 1963) *horse lotteries* defined on a denumerable reward set $\{\boldsymbol{r}_1, \boldsymbol{r}_2, \ldots\}$. A (simple) *horse lottery* h is a function from states to (simple) probability distributions, i.e., von Neumann–Morgenstern lotteries over the set of rewards. Denote such a lottery by L. Then $h(\omega_j) = L_j$. In von Neumann–Morgenstern theory, with $0 \leq x \leq 1$, the x: $(1-x)$ convex combination of two lotteries, L_1 and L_2, denoted $L_3 = xL_1 \oplus (1-x)L_2$, is the (simple) probability distribution obtained by the x: $(1-x)$ weighted average of the two (simple) distributions that define L_1 and L_2. Following Anscombe–Aumann's theory, with $0 \leq x \leq 1$, define the convex combination of two horse lotteries, $h_3 = xh_1 \oplus (1-x)h_2$, by $h_3(\omega_j) = L_{3j} = xL_{1j} \oplus (1-x)L_{2j}$.

We use the topology of pointwise convergence in distributions to define closure of a set of options. That is, a sequence of horse lotteries $< h_i : i = 1, \ldots >$ converges to the horse lottery h if, for each $1 \leq j \leq n$, the sequence $< L_{ij} : i = 1, \ldots >$ converges to the lottery L_j.

Let $\boldsymbol{H}(\boldsymbol{O})$ denote the result of taking the (closed) convex hull of the option set $\boldsymbol{O}$. That is, $\boldsymbol{H}(\boldsymbol{O})$ is the (closure of the) set of all mixed options in $\boldsymbol{O}$. Then when $\boldsymbol{O}$ is finite, without loss of generality, q^* of Theorem 1 may be taken to be an option that also is Bayes-admissible for some $p^* \in \boldsymbol{P}$. That is, in Theorem 1 we may choose $q^* \in \boldsymbol{H}(\boldsymbol{O})$ such that $q^* \in \boldsymbol{C}(\boldsymbol{H}(\boldsymbol{O}))$ for a coherent choice function using the set $\boldsymbol{P}$ of all probability distributions on Ω. In terms of Theorem 1, in Example 1 with $o^* = h$, then $q_x^* = \mathrm{x} f \oplus (1 - \mathrm{x})\, g$ for $.4 < x < .6$ uniformly dominates o^*. But each such q_x^* is Bayes with respect to $\boldsymbol{H}(\boldsymbol{O})$ precisely for one probability on $\Omega : \boldsymbol{p}\,(\omega_1) = .5$. We use this fact, next, to establish that each set of probabilities has its own *unique* coherent choice function.

Aside: Under a topology of pointwise convergence in lotteries, Theorem 1 generalizes to infinite states spaces Ω and infinite, closed options sets $\boldsymbol{O}$ by using Theorem 2.1 of Kindler (1983) to replace Pearce's use of von Neumann's Minimax Theorem, which does not generalize to infinite games. But then the mixed strategies needed with Kindler's result are merely finitely additive, rather than countably additive. For an intermediate generalization using countably additive mixed strategies, where the set of feasible options may be infinite though the state space is finite, see our (2008).

In this section we use decision problems involving horse lotteries defined on only two privileged rewards, **0** and **1**, to individuate different sets of indeterminate probabilities. Regarding the two privileged rewards, we assume there is a strict preference for the constant horse lottery **1** over the constant horse lottery **0**. That is, $\boldsymbol{C}\{\mathbf{0}, \mathbf{1}\} = \{\mathbf{1}\}$. We consider coherent choice using state-independent utilities, where for each utility, $\boldsymbol{u}(\mathbf{1}) = 1$ and $\boldsymbol{u}(\mathbf{0}) = 0$ in each state, ω. Since in this construction we use horse lotteries involving only these two rewards, our goal is to show that if $\boldsymbol{P}$ and $\boldsymbol{P}'$ are two different sets of probabilities, the coherent choice function based on the set $\boldsymbol{S} = \boldsymbol{P} \times \{\boldsymbol{u}\}$ is different from the coherent choice function based on the set $\boldsymbol{S}' = \boldsymbol{P}' \times \{\boldsymbol{u}\}$. We establish that goal as a corollary to the Theorem 2, below, which shows how to use coherent choice functions to characterize membership of a particular distribution p in an arbitrary set of distributions, $\boldsymbol{P}$.

Let $\boldsymbol{p}^* = (p_1, \ldots, p_n)$ be a probability distribution on Ω. Denote by $\underline{p}$ the smallest nonzero coordinate of $\boldsymbol{p}^*$. Define the *constant* horse lottery act $\boldsymbol{a} = \underline{p}\bar{\mathbf{1}} + (1 - \underline{p})\mathbf{0}$, which yields the same lottery in each state. For each $j = 1, \ldots, n$, *define* the act h_j as follows.

$$h_j(\omega_i) = \begin{cases} 1 & \text{if } \; i = j \; \text{ and } \; p_j = 0 \\ a & \text{if } \; i \neq j \; \text{ and } \; p_j = 0 \\ \frac{\underline{p}}{p_j} 1 \oplus \left(1 - \frac{\underline{p}}{p_j}\right) 0 & \text{if } \; i = j \; \text{ and } \; p_j > 0 \\ 0 & \text{if } \; i \neq j \; \text{ and } \; p_j > 0 \end{cases}$$

Consider the finite option set $\boldsymbol{O}_{\boldsymbol{p}^*} = \{\boldsymbol{a}, h_1, \ldots, h_{\mathrm{n}}\}$. Let $\boldsymbol{P}$ be a non-empty set of probabilities defined on Ω and let $\boldsymbol{C}(\cdot)$ be a coherent choice function based on the non-empty set $\boldsymbol{S}$ of probability-utility pairs of the form $\boldsymbol{S} = \boldsymbol{P} \times \{\boldsymbol{u}\}$. Denote by $\mathbf{E}_{\boldsymbol{p}}(\boldsymbol{u}(\cdot))$ the expected utility function with respect to a pair $(\boldsymbol{p}, \boldsymbol{u})$ in $\boldsymbol{S}$.

Theorem 2 $\boldsymbol{p}^* \in \boldsymbol{P}$ *if and only if* $\boldsymbol{C}(\boldsymbol{O}_{\boldsymbol{p}^*}) = \boldsymbol{O}_{\boldsymbol{p}^*}$.

Proof First, observe that for all j and utility $\boldsymbol{u}$, $\mathbf{E}_{\boldsymbol{p}^*}(\boldsymbol{u}(h_{\mathrm{j}})) = \underline{p} = \mathbf{E}_{\boldsymbol{p}^*}(\boldsymbol{u}(\boldsymbol{a}))$. For the "only if" direction, assume that the pair $(\boldsymbol{p}^*; \boldsymbol{u}) \in \boldsymbol{S}$ for some utility $\boldsymbol{u}$. Then by this equality, every element of $\boldsymbol{O}_{\boldsymbol{p}^*}$ is Bayes with respect to $(\boldsymbol{p}^*; \boldsymbol{u})$ and $\mathrm{C}(\boldsymbol{O}^*_{\boldsymbol{p}}) = \boldsymbol{O}_{\boldsymbol{p}^*}$. For the "if" direction, assume that $\boldsymbol{C}(\boldsymbol{O})^*_{\boldsymbol{p}} = \boldsymbol{O}_{\boldsymbol{p}^*}$. Note that $\mathbf{E}_{\mathrm{q}}(\boldsymbol{u}(\boldsymbol{a})) = \underline{p}$ for every probability/utility pair $(\boldsymbol{q}, \boldsymbol{u})$. Let $(\boldsymbol{q}, \boldsymbol{u})$ be a probability/utility pair with $\boldsymbol{q} \neq \boldsymbol{p}^*$. First, consider the case with $\underline{p} < 1$. Then there exists j with $\boldsymbol{q}_j > \boldsymbol{p}_j$. So,

$$\mathbf{E}_{\mathrm{q}}(\boldsymbol{u}(h_{\mathrm{j}})) = \begin{cases} q_{\mathrm{j}}\underline{p}/p_{\mathrm{j}} > \underline{p} & \text{if} \quad p_{\mathrm{j}} > 0, \\ q_{\mathrm{j}} + (1 - q_{\mathrm{j}})\underline{p} > \underline{p} & \text{if} \quad p_{\mathrm{j}} = 0. \end{cases}$$

Hence, for each $(\boldsymbol{q}, \ \boldsymbol{u})$ with $\boldsymbol{q} \neq \boldsymbol{p}^*$, $\mathbf{E}_{\mathbf{q}}(\boldsymbol{u}(h_{\mathrm{j}})) > \mathbf{E}_{\mathbf{q}}(\boldsymbol{u}(\boldsymbol{a}))$. It follows that $\boldsymbol{a} \notin \boldsymbol{C}(\boldsymbol{O}_{\boldsymbol{p}^*})$ unless $(\boldsymbol{p}^*, \boldsymbol{u}) \in \boldsymbol{S}$. Finally, consider the case with $\underline{p} = 1$. In this case, $\boldsymbol{O}_{\boldsymbol{p}^*} = \{\mathbf{1}, h_j\}$ where $p_{\mathrm{j}} = 1$. So, $\mathbf{E}_{\mathbf{q}}(\boldsymbol{u}(h_{\mathrm{j}})) = q_{\mathrm{j}} < 1 = \mathbf{E}_{\mathbf{q}}(\boldsymbol{u}(\boldsymbol{a}))$ for every probability/utility pair $(\boldsymbol{q}, \boldsymbol{u})$ with $\boldsymbol{q} \neq \boldsymbol{p}^*$. It follows that $h_j \notin \boldsymbol{C}(\boldsymbol{O}_{\boldsymbol{p}^*})$ unless $(\boldsymbol{p}, \ \boldsymbol{u}) \notin \boldsymbol{S}$. □

Corollary *Let* $\boldsymbol{P}_1$ *and* $\boldsymbol{P}_2$ *be two distinct (nonempty) sets of probabilities with corresponding Bayes-admissible choice functions* $\boldsymbol{C}_1$ *and* $\boldsymbol{C}_2$*. There exists a finite option set* $\boldsymbol{O}_{\boldsymbol{p}}$*, as above, such that* $\boldsymbol{C}_1(\boldsymbol{O}_p) \neq \boldsymbol{C}_2(\boldsymbol{O}_{\boldsymbol{p}})$.

Proof Since $\boldsymbol{P}_1 \neq \boldsymbol{P}_2$, either there exists $\boldsymbol{p} \in \boldsymbol{P}_1$ and $\boldsymbol{p} \notin \boldsymbol{P}_2$ or, $\boldsymbol{p} \notin \boldsymbol{P}_1$ and $\boldsymbol{p} \in \boldsymbol{P}_2$. Construct the finite option set $\boldsymbol{O}_{\boldsymbol{p}}$ as above. Then by Theorem 2, $\boldsymbol{C}_1(\boldsymbol{O}_{\boldsymbol{p}}) \neq \boldsymbol{C}_2(\boldsymbol{O}_{\boldsymbol{p}})$. □

Thus, each set of probabilities $\boldsymbol{P}$ has its own distinct pattern of Bayes-admissible choice functions with respect to option sets $\mathbf{O}_{\boldsymbol{p}}$ for $\boldsymbol{p} \in \boldsymbol{P}$.

> *Aside*: This Corollary is a generalization of Theorem 1 from our (2004) paper, which establishes distinct coherent choice functions for distinct *convex* sets $\boldsymbol{p}$ of probabilities.

3 Axiomatizing coherent choice functions

We turn, next, to a system of axioms for choice functions that are necessary for coherence, and which are jointly sufficient for a representation of choice by a set $\boldsymbol{S}$ of probability/almost-state-independent utility pairs, as explained below. We provide sufficient conditions when these pairs have a common state-independent utility. In such a case the coherent choice function corresponds to choice under indeterminate uncertainty with a determinate utility.

We continue with the framework of the previous section: options are simple horse lotteries defined over a finite state space $\Omega = \{\omega_1, \ldots, \omega_n\}$. Using choice functions over sets of options permits us to extend our (1995) work, which deals solely with (partially ordered) strict preference $\prec$. Specifically, interpret the strict preference relation, $h_1 \prec h_2$ as fixing coherent choice in binary option sets: $\boldsymbol{C}\{h_1, h_2\} = \{h_2\}$.

Then, as explained below, our (1995) theory of strict partial orders is a special case of coherent choice functions. Thus, some results that follow from binary choice problems are available also within this theory. For example, it then follows from Sect. II.6 of our (1995) theory that each agreeing cardinal utility for the choice function $\boldsymbol{C}(\cdot)$, if one exists, is a bounded utility function, since that condition is already forced by considering choice problems with binary option sets.

Also, with this interpretation of our (1995) theory of binary choice, we assert without proof the following characterization of the Maximality rule. Specifically, let $\boldsymbol{C}(\cdot)$ be a coherent choice function based on a set $\boldsymbol{S}$ of probability/utility pairs.

Definition Define the choice function $\boldsymbol{C}_M(\cdot)$ using $\boldsymbol{C}(\cdot)$ applied to pairs of options in $\boldsymbol{O}$ as follows. Declare an option o in $\boldsymbol{O}$ to be $\boldsymbol{C}_M$-admissible *if and only if* it is $\boldsymbol{C}$-admissible in all binary problems, i.e., $o \in \boldsymbol{C}(\{o, o'\})$ for each $\left\{o, o'\right\} \subseteq \boldsymbol{O}$.

Proposition *A choice function accords with the Maximality rule of admissibility for the set* $\mathbf{S}$ *if and only if it is of the form* $\boldsymbol{C}_M(\cdot)$ *for the coherent choice function* $\boldsymbol{C}(\cdot)$ *based on* $\mathbf{S}$.

Here, we focus on an axiomatic representation of coherent choice when utility is determinate regarding the two distinguished prizes **1** and **0**, which we take, respectively, as the upper and lower bounds on all other constant acts: the constant act **1** is better than, and the constant act **0** is worse than, each other constant act. For convenience, we assume that all cardinal utilities are scaled so that $\mathbf{u}(\mathbf{1}) = 1$ and $\boldsymbol{u}(\mathbf{0}) = 0$.

With these assumptions about the utility for the two distinguished rewards, and to motivate our axiomatization of coherent choice functions, next we rehearse a standard axiomatization of the Anscombe–Aumann (1963) theory of binary preference ($\preceq$).

A-A Axiom 1 Choice over sets of horse lotteries reduces to a pairwise comparison of options by preference ($\preceq$), which generates an *Ordering* of the set of options: $\preceq$ is reflexive, transitive, and complete for all pairs. That is, there exists a binary preference order $\preceq$ over $\mathcal{H} \times \mathcal{H}$ such that an option h is admissible from a feasible set $\boldsymbol{O}$ if and only if h is $\preceq$-maximal in $\boldsymbol{O}$.

Definition Let $\prec$ denote *strict preference*, the asymmetric part of $\preceq$: $h_1 \prec h_2$ *if and only if* $h_1 \preceq h_2$ and not $h_2 \preceq h_1$; and let $\approx$ denote *indifference*, the symmetric part of $\preceq$: $h_1 \approx h_2$ *if and only if* $h_1 \preceq h_2$ and $h_2 \preceq h_1$.

A-A Axiom 2 Preference ($\preceq$) satisfies the von Neumann–Morgenstern postulate of *Independence*. For each h_1, h_2 and h_3, and for each $0 < x \leq 1$, $h_1 \prec h_2$ *if and only if* $xh_1 \oplus (1-x)h_3 \prec xh_2 \oplus (1-x)h_3$

A-A Axiom 3 An Archimedean axiom—to secure that preference ($\preceq$) admits a real-valued representation, thus insuring also a real-valued representation for subjective probability over Ω and a real-valued cardinal utility over prizes.

If $h_1 \prec h_2 \prec h_3$ there exist $0 < x,\ y < 1$ such that $xh_1 \oplus (1-x)h_3 \prec h_2 \prec xh_1 \oplus (1-x)h_3$

A-A Axiom 4 For existence of a state-independent utility representation for preference, a final axiom requires that the decision maker's preference for constant horse lotteries reproduces under each non-null state in the form of called-off horse lotteries. This is made precise as follows.

Definition A state $\omega^* \in \Omega$ is *null* provided that for all pairs of horse lotteries, $h_1 \approx h_2$ whenever $h_1(\omega) = h_2(\omega)$ for $\omega \neq \omega^*$. It is a *non-null* state otherwise.

A pair of horse lotteries h_1, h_2 are *called-off* on the event E if for each $\omega \in E, h_1(\omega) = h_2(\omega)$.

Let h_1 and h_2 be a pair of constant horse lotteries yielding, respectively the reward r_1 and r_2. That is, for each $\omega \in \Omega$, $h_1(\omega) = r_1$ and $h_2(\omega) = r_2$. Let ω^* be a non-null state. Axiom 4 requires that, for each pair of horse lotteries h_{1^*} and h_{2^*} such that $h_{1^*}(\omega^*) = r_1$ and $h_{2^*}(\omega^*) = r_2$, and they are called-off on $\{\omega^*\}^{\mathbf{c}}$, then $h_1 \preceq h_2$ *if and only if* $h_{1^*} \preceq h_{2^*}$

We adapt our presentation here to match these four axioms. For ease of exposition some conditions are formulated in terms of the *rejection function*, $\boldsymbol{R}(\cdot)$ which identifies the $\boldsymbol{C}$-inadmissible options from a feasible set $\boldsymbol{O}$.

Definition $\boldsymbol{R}(\boldsymbol{O}) = \boldsymbol{O} - \boldsymbol{C}(\boldsymbol{O})$

In place of the ordering axiom, we require the following two conditions:

Axiom 1a—Sen's property *alpha* If $\boldsymbol{O}_2 \subseteq \boldsymbol{R}(\boldsymbol{O}_1)$ and $\boldsymbol{O}_1 \subseteq \boldsymbol{O}_3$, then $\boldsymbol{O}_2 \subseteq \boldsymbol{R}(\boldsymbol{O}_3)$.

You cannot promote an unacceptable option into an acceptable option by adding options to the feasible set.

Axiom 1b—a variant of Aizerman's (1985) condition, if $\boldsymbol{O}_2 \subseteq \boldsymbol{R}(\boldsymbol{O}_1)$ and $\boldsymbol{O}_3 \subseteq \boldsymbol{O}_2$, then $\boldsymbol{O}_2 - \boldsymbol{O}_3 \subseteq \boldsymbol{R}(closure[\boldsymbol{O}_1 \subseteq \boldsymbol{O}_3])$.

You cannot promote an unacceptable option into an acceptable option by deleting unacceptable options from the option set.

Note: We require *closure of* $[\boldsymbol{O}_1 - \boldsymbol{O}_3]$ since $\boldsymbol{O}_1 - \boldsymbol{O}_3$ may not be a closed set, despite the fact that $\boldsymbol{O}_1$ and $\boldsymbol{O}_3$ are closed.

With Axioms 1a and 1b, define a strict partial order $\langle$ on sets of options as follows. Let $\boldsymbol{O}_1$ and $\boldsymbol{O}_2$ be two option sets.

Definition $\boldsymbol{O}_1 \langle \boldsymbol{O}_2$ if and only if $\boldsymbol{O}_1 \subseteq \boldsymbol{R}[\boldsymbol{O}_1 \cup \boldsymbol{O}_2]$.
So $\boldsymbol{O}_1 \langle \boldsymbol{O}_2$ obtains when $\boldsymbol{O}_1$ contains only inadmissible options in a choice among the options in both sets, $\boldsymbol{O}_1 \cup \boldsymbol{O}_2$. Lemma 1 of our (2004) establishes that given Axioms 1a and 1b, the binary relation $\langle$ is a strict partial order over pairs of sets of options: $\langle$ is transitive and irreflexive. This finding is the basis for the assertion, above, that our (1995) theory of strict partial orders is a special case of the theory developed here for coherent choice functions. Our (1995) theory is the restriction of the current theory to decision problems with two feasible options.

The role of mixtures between options is captured in the following pair of axioms for $\langle$. With $\boldsymbol{O}_1$ an option set and o an option, the notation $\alpha\boldsymbol{O}_1 \oplus (1-\alpha)o$ denotes the set of pointwise mixtures, $\alpha o_1 \oplus (1-\alpha)o$ for $o_1 \in \boldsymbol{O}_1$.

Axiom 2a—*Independence* is formulated for the relation $\langle$ over sets of options. Specifically, let o be an option and $0 < \alpha \leq 1$.

$\boldsymbol{O}_1 \langle \boldsymbol{O}_1$ if and only if $\alpha \boldsymbol{O}_1 \oplus (1-\alpha) o \langle \alpha \boldsymbol{O}_2 \oplus (1-\alpha) o$.

Axiom 2b—*Mixtures* If $o \in \boldsymbol{O}$ and $o \in \boldsymbol{R}[\boldsymbol{H}(\boldsymbol{O})]$, then $o \in \boldsymbol{R}[\boldsymbol{O}]$.

Axiom 2b asserts that inadmissible options from a mixed set remain so even before mixing.

With respect to the three decision rules discussed in Sect. 1, *Independence* (Axiom 2a) fails in Γ-*Maximin* theory. *Mixing* (Axiom 2b) fails for the choice function determined by *Maximality*. Thus, by distinguishing between axioms 2a and 2b we highlight what we judge to be the key difference between Γ-*Maximin* and *Maximality*, and how each fails to limit admissibility to Bayes-admissibility.

The Archimedean condition for coherent choice functions requires a technical adjustment from the canonical form used by, e.g. von Neumann–Morgenstern theory or Anscombe–Aumann theory. The canonical form is too restrictive in this setting. (See Sect. II.4 of our 1995.) The reformulated version of the Archimedean condition is as a continuity principle compatible with strict preference as a strict partial order. It reads as follows.

Let $\boldsymbol{A}_{\mathrm{n}}$ and $\boldsymbol{B}_{\mathrm{n}} (n = 1, \ldots)$ be sets of options converging pointwise, respectively, to the option sets $\boldsymbol{A}$ and $\boldsymbol{B}$. Let $\boldsymbol{N}$ be an option set.

Axiom 3a If, for each n, $\boldsymbol{B}_{\mathrm{n}} \langle \boldsymbol{A}_{\mathrm{n}}$ and $\boldsymbol{A} \langle \boldsymbol{N}$, then $\boldsymbol{B} \langle \boldsymbol{N}$.
Axiom 3b If, for each n, $\boldsymbol{B}_{\mathrm{n}} \langle \boldsymbol{A}_{\mathrm{n}}$ and $\boldsymbol{N} \langle \boldsymbol{B}$, then $\boldsymbol{N} \langle \boldsymbol{A}$.

The counterpart to A-A Axiom 4 for state-neutrality is captured by the following dominance relations. Consider horse lotteries $\boldsymbol{h}_1$ and $\boldsymbol{h}_2$, with $\boldsymbol{h}_i(\omega_j) = \beta_{ij}\mathbf{1} \oplus (\mathbf{1} - \beta_{ij})\mathbf{0}$; $i = 1, 2$ $j = 1, \ldots, n$.

Definition $\boldsymbol{h}_2$ ***weakly dominates*** $\boldsymbol{h}_1$ if $\beta_{2j} \geq \beta_{1j}$ for $j = 1, \ldots, n$.

Assume that o_2 weakly dominates o_1, and that a is an option different from each of these two.

Axiom 4a If $o_2 \in \boldsymbol{O}$ and $a \in \boldsymbol{R}(\{o_1\} \cup \boldsymbol{O})$ then $a \in \boldsymbol{R}(\boldsymbol{O})$.
Axiom 4b If $o_1 \in \boldsymbol{O}$ and $a \in \boldsymbol{R}(\boldsymbol{O})$ then $a \in \boldsymbol{R}([\{o_2\} \cup \boldsymbol{O} - \{o_1\}])$.

In words, Axiom 4*a* says that when a weakly dominated option is removed from the set of options, other inadmissible options remain inadmissible. So, by Axiom 1, when an option is replaced in the option set by one that it weakly dominates, other admissible options remain admissible.

Axiom 4*b* says that when an option is replaced by one that weakly dominates it, (other) inadmissible options remain inadmissible. Trivially by Axiom 1, merely adding a weakly dominating option cannot promote an inadmissible option into one that is admissible.

Axiom 4 captures key aspects of what Savage's postulate **P3** asserts about state-independent utility of the prizes **1** and **0** without assuming states are not-null. That is, the intended representation for a coherent choice function $\boldsymbol{C}(\cdot)$ uses the expected utility rule (i.e., Bayes-admissibility) applied with a set $\boldsymbol{S}$ of probability/utility pairs. However, it may be that for each state ω_j there is a probability/utility pair, $(\boldsymbol{p}_{\mathbf{j}}, \boldsymbol{u}) \in \boldsymbol{S}$

such that $\boldsymbol{p_j}(\omega_j) = 0$. In the language of our (1995) paper, then each state in Ω is *potentially* null under $\boldsymbol{S}$. Thus, Savage's **P3** (or the corresponding Anscombe–Aumann Axiom 4) is vacuous when potentially null states are excepted. Nonetheless, Axiom 4 reports two facts about weakly dominated lotteries that obtain even when each state is potentially null.

Theorem 3 *Axioms* 1–4 *are necessary for a coherent choice function.*

Let $\boldsymbol{S}$ *be a non-empty set of pairs of probability/state-independent utilities, and let* $\boldsymbol{C_S}(\cdot)$ *be the coherent choice function defined by setting the admissible options in feasible set* $\boldsymbol{O}$ *to be exactly those that are Bayes-admissible with respect to* $\boldsymbol{S}$. *Then* $\boldsymbol{C_S}(\cdot)$ *satisfies Axioms* 1–4.

Proof The argument for the necessity of Axioms 1–3 is given in our (2004). That Axiom 4 is necessary as well follows immediately by noting that whenever o_2 weakly dominates o_1 then for each $(\boldsymbol{p},\boldsymbol{u}) \in \boldsymbol{S}$, $\mathbf{E}_{\boldsymbol{p}}(\boldsymbol{u}(o_2)) \geq \mathbf{E}_{\boldsymbol{p}}(\boldsymbol{u}(o_1))$. □

The following result is helpful in linking our theory with Theorem 1. Consider horse lotteries $\boldsymbol{h}_1$ and $\boldsymbol{h}_2$, with $\boldsymbol{h}_i(\omega_j) = \beta_{ij}\mathbf{1} \oplus (1-\beta_{ij})\mathbf{0}$; $i = 1, 2$ $j = 1, \ldots, n$.

Definition $\boldsymbol{h}_2$ ***strongly dominates*** $\boldsymbol{h}_1$ if $\beta_{2j} > \beta_{1j}$ for $j = 1, \ldots, n$.

Lemma 1 ***Inadmissibility of strongly dominated options***: *If* h_2 *strongly dominates* h_1 *then* $\{h_1\} = \boldsymbol{R}(\{h_1, h_2\})$.

Proof The strategy of the proof is as follows: Use the Independence axiom to convert the problem with option set $\boldsymbol{O} = \{h_1, h_2\}$ into an equivalent problem $\boldsymbol{O}' = \{h_1', h_2'\}$, where h_1' is a constant horse lottery, and where h_2' strongly dominates h_1'. Then we show that h_2' weakly dominates another constant horse lottery, h_2'' which also strongly dominates h_1'. Then, by Independence $\{h_1'\} = \boldsymbol{R}(\{h_1', h_2''\})$ and by Axiom 4*b*, $\{h_1'\} = \boldsymbol{R}(\{h_1', h_2'\})$. Last, by Independence, $\{h_1\} = \boldsymbol{R}(\{h_1, h_2\})$.

Here are the details. Let $0 \leq \beta* = min\{\beta_{1j}\}$ and $1 > \beta^* = max\{\beta_{1j}\}$. Let $h_3(\omega_j) = \beta_{3j}\mathbf{1} \oplus (1-\beta_{3j})\mathbf{0}$, where $\beta_{3j} = \beta^* + \beta* - \beta_{1j}$. Then the horse lottery $h_\mathbf{1}' = .5h_\mathbf{1} \oplus .5h_\mathbf{3}$ is the constant (von Neumann–Morgenstern) lottery with $\beta_{1j}' = (\beta^* + \beta*)/2$. Define $h_\mathbf{2}' = .5h_\mathbf{2} \oplus .5h_\mathbf{3}$. The Independence axiom asserts that $\{h_1\} = \boldsymbol{R}(\{h_1, h_2\})$ if and only if $\{h_1'\} = \boldsymbol{R}(\{h_1', h_2'\})$. But h_2' strongly dominates h_1', because h_2 strongly dominates h_1. In fact, $\beta_{2j}' - \beta_{1j}' = (\beta_{2j} - \beta_{1j})/2 > 0$. So, let $0 < \delta = min\{\beta_{2j} - \beta_{1j}\}$, and then $\delta/2 = min\{\beta_{2j}' - \beta_{1j}'\}$. Let h_2'' be the constant (von Neumann–Morgenstern) lottery defined with $\beta_{2j}'' = \beta_{1j}' + \delta/2 = (\beta^* + \beta* + \delta)/2 > \beta_{1j}'$. Observe, also, that h_2' weakly dominates h_2''. Then, as announced before, by Independence $\{h_1'\} = \boldsymbol{R}(\{h_1', h_2''\})$; by Axiom 4*b*, $\{h_1'\} = \boldsymbol{R}(\{h_1', h_2'\})$; and by another application of Independence, $\{h_1\} = \boldsymbol{R}(\{h_1, h_2\})$. □

Next we introduce two concepts central to our argument for representing coherent choice functions.

Definition The pair $(\boldsymbol{p}, \boldsymbol{u})$ is *a local Bayes model for option* $\boldsymbol{o}$ provided that $\boldsymbol{o}$ maximizes $(\boldsymbol{p}, \boldsymbol{u})$-expected utility with respect to the options in set $\boldsymbol{O}$.

The pair (***p,u***) is it a <u>*global Bayes model for the choice function*</u> ***C***(•) provided that, for each option set ***O***, if $o \in \boldsymbol{O}$ maximizes (***p,u***)-expected utility with respect to the options in set ***O*** then $o \in \boldsymbol{C}(\boldsymbol{O})$.

We adapt the concept of a set of *almost state-independent utilities*, presented in our (1995, Definition 31), as follows. Let $\{\boldsymbol{r}_1, \ldots, \boldsymbol{r}_m\}$ be a set of rewards and assume that for each constant horse lottery $\boldsymbol{r} \in \{\boldsymbol{r}_1, \ldots, \boldsymbol{r}_m\}\{\mathbf{0}\}\langle\{r\}\langle\{\mathbf{1}\}$, so that the constant acts **0** and **1** strictly bound the value of the other constant acts.

The set of probability/utility pairs $\boldsymbol{S}^{\#} = \{(\boldsymbol{p}_j, \boldsymbol{u}_j) : j = 1, \ldots\}$ form a set of <u>*almost state independent utilities*</u> for $\{\boldsymbol{r}_1, \ldots, \boldsymbol{r}_m\}$ provided that for each $\varepsilon > 0$, there is a pair $(\boldsymbol{p}_\varepsilon, \boldsymbol{u}_\varepsilon) \in \boldsymbol{S}^{\#}$ and a set of states $\Omega_{(1-\varepsilon)} \subseteq \Omega$ with $\boldsymbol{p}_\varepsilon(\Omega_{(1-\varepsilon)}) \geq 1 - \varepsilon$ such that for each $r\varepsilon\{\boldsymbol{r}_1, \ldots, \boldsymbol{r}_m\}$.

$$\max_{\omega_i, \omega_j \in \Omega_{1-\varepsilon}} | u_{\varepsilon, \omega_i}(r) - u_{\varepsilon, \omega_j}(r)| \leq \varepsilon.$$

This allows us to formulate the central result of this paper.

Theorem 4 *A choice function* ***C***(·) *defined on the class* $\mathcal{H}$ *of simple Anscombe–Aumann Horse-lotteries using (at least) three prizes* $\{\mathbf{0}, \boldsymbol{r}, \mathbf{1}\}$, *with* $\{\mathbf{0}\}\langle\{\boldsymbol{r}\}\langle\{\mathbf{1}\}$, *satisfies our 4 axioms only if it is represented as a coherent choice function by a set* ***S*** *of probability/almost-state-independent utility pairs.*

> *Aside*: A sufficient condition is given at the end of Appendix 2 for the global Bayes models of ***S*** to be comprised solely of probability/state-independent utility pairs.

Proof This theorem follows from three lemmas.

Lemma 2 *For each choice set* ***O*** *and admissible option* $o \in \boldsymbol{C}(\boldsymbol{O})$, *o has at least one local Bayes model.*

Proof By Theorem 1, an option lacking a local Bayes model is strongly dominated by a finite mixture of other options already available in the same choice problem. Then, Axiom 3 and the Lemma 1 on inadmissibility of strongly dominated options demonstrate Lemma 2. □

> *Aside*: Let $o \in \boldsymbol{C}(\boldsymbol{O})$. If (***p,u***) and $(\boldsymbol{p}', \boldsymbol{u})$ both are local Bayes models for *o*, then so too is each pair (***q, u***) of the form $\boldsymbol{q} = x\boldsymbol{p} + (1 - x)\boldsymbol{p}'(0 \leq x \leq 1)$. Likewise, if each of $(\boldsymbol{p}_j, \boldsymbol{u})(j = 1, \ldots)$ is a local Bayes model for *o* and the sequences of distributions $\{\boldsymbol{p}_j\}$ converges to distribution ***q***, then also (***q, u***) is a local Bayes model for *o*. Hence, we have the following corollary

Corollary *The set of local Bayes models for* $o \in \boldsymbol{C}(\boldsymbol{O})$ *with a common utility* ***u*** *is given by a non-empty, closed, convex set of probabilities.*

Let $\mathcal{H}_{\{\mathbf{0},\mathbf{1}\}}$ be the set of horse lotteries defined using only the privileged pair of rewards, 0 and 1. Based on the ideas presented in Sect. 2, given a distribution ***p***, next we define a special choice problem ***O**** involving only horse lotteries from $\mathcal{H}_{\{\mathbf{0},\mathbf{1}\}}$ so that, precisely when all of ***O****'s options are admissible, then ***p*** is a global Bayes model for the choice function defined on feasible subsets of $\mathcal{H}_{\{\mathbf{0},\mathbf{1}\}}$. Thus, the notation for the special choice problem should be '$\boldsymbol{O}^*_{\boldsymbol{p}}$' with the subscript identifying the distribution $\boldsymbol{p}$. To make the proofs readable, that subscript is suppressed here.

Lemma 3 *Suppose that* $\boldsymbol{C}(\boldsymbol{O}^*) = \boldsymbol{O}^*$. *Then* $\boldsymbol{p}$ *is a global Bayes model for the choice function* $\boldsymbol{C}(\cdot)$ *restricted to feasible sets in* $\mathcal{H}_{\{\mathbf{0},\mathbf{1}\}}$.

Proof See Appendix 1.

Lemma 4 *For each admissible option* $o \in \boldsymbol{C}(\boldsymbol{O})$ *at least one of its local Bayes models is a global Bayes model or else there is a set of probability/almost-state-independent utility pairs that represent* $\boldsymbol{C}$.

Proof See Appendix 2.

4 An example of coherent choice using a non-convex set $\boldsymbol{P}$ reflecting "expert" opinion

In this section we illustrate how coherent choices may represent "expert" opinions while preserving independence between two events. The following example highlights the use of a non-convex set of probabilities to represent a coherent choice function.

Example 2 Consider a decision problem among three options—three treatment plans $\{T_1, T_2, T_3\}$ defined over four states $\Omega = \{\omega_1, \omega_2, \omega_3, \omega_4\}$ with determinate utility outcomes given in the following table. That is, the numbers in Table 1 are the utility outcomes for the options (rows) in the respective states (columns).

Let a convex set $\boldsymbol{P}$ of probabilities be generated by two extreme points, distributions $\boldsymbol{p}_1$ and $\boldsymbol{p}_2$, defined in Table 2. Distribution $\boldsymbol{p}_3$ is the .50–.50 (convex) mixture of $\boldsymbol{p}_1$ and $\boldsymbol{p}_2$.

Note that (for $i = 1, 2, 3$) under probability $\boldsymbol{p}_i$, only option T_i is Bayes-admissible from the option set of $\{T_1, T_2, T_3\}$. Without convexity—that is, using the set $\boldsymbol{P}$ containing the two (extreme) distributions $\boldsymbol{P} = \{\boldsymbol{p}_1, \boldsymbol{p}_2\}$—option T_3 is the sole Bayes-inadmissible option from among the three options $\{T_1, T_2, T_3\}$.

Now, interpret these states as the cross product of two binary partitions: a medical event $\boldsymbol{A}$ (patient allergic) and its complementary event $\boldsymbol{NA}$ (patient not-allergic), with a binary meteorological partition. $\boldsymbol{S}$ (sunny) and $\boldsymbol{NS}$ (cloudy). Specifically:

Table 1 Utilities for each of the three treatment plans in each of the four states

	ω_1	ω_2	ω_3	ω_4
T_1	0.00	0.00	1.00	1.00
T_2	1.00	1.00	0.00	0.00
T_3	0.99	−0.01	−0.01	0.99

Table 2 Probabilities for the four states under each of the three distributions

	ω_1	ω_2	ω_3	ω_4
$\boldsymbol{p}_1$	0.08	0.32	0.12	0.48
$\boldsymbol{p}_2$	0.48	0.12	0.32	0.08
$\boldsymbol{p}_3$	0.28	0.22	0.22	0.28

$\omega_1 = \boldsymbol{A}\&\boldsymbol{S}$, $\omega_2 = \boldsymbol{A}\&\boldsymbol{NS}$, $\omega_3 = \boldsymbol{NA}\&\boldsymbol{S}$, $\omega_4 = \boldsymbol{NA}\&\boldsymbol{NS}$. Under probability distribution $\boldsymbol{p}_1$, the two partitions are independent events with $\boldsymbol{p}_1(\boldsymbol{A}) = .4$ and $\boldsymbol{p}_1(\boldsymbol{S}) = .2$. Likewise, under probability distribution $\boldsymbol{p}_2$, the two partitions are independent events with $\boldsymbol{p}_2(\boldsymbol{A}) = .6$ and $\boldsymbol{p}_2(\boldsymbol{S}) = .8$. And under distribution $\boldsymbol{p}_3$ the events $\boldsymbol{A}$ and $\boldsymbol{S}$ are positively correlated: $.56 = \boldsymbol{p}_3(\boldsymbol{A}|\boldsymbol{S}) > \boldsymbol{p}_3(\boldsymbol{A}) = .5$, as happens with each distribution $\boldsymbol{q}$ that is a non-trivial mixture of $\boldsymbol{p}_1$ and $\boldsymbol{p}_2$.

Then the three options have the following interpretations: T_1 and T_2 are ordinary medical options for how to treat the patient, with outcomes that depend solely upon the patient's allergic state. T_3 is an option that makes the allocation of medical treatment a function of the meteorological state, with a "fee" of 0.01 utile assessed for that input. That is, T_3 is the option "T_1 if cloudy and T_2 if sunny, while paying a fee of 0.01."

Suppose that $\boldsymbol{p}_1$ represents the opinion of medical expert 1, and $\boldsymbol{p}_2$ represents the opinion of medical expert 2. Without convexity of the credal probabilities, T_3 is inadmissible. This captures the shared agreement between the two medical experts that T_3 is unacceptable from the choice of three options $\{\mathrm{T}_1, \mathrm{T}_2, \mathrm{T}_3\}$, and it captures the pre-systematic understanding that under T_3 you pay to use medically irrelevant inputs about the weather in order to determine the medical treatment. However, with convexity of the set generated by $\boldsymbol{p}_1$ and $\boldsymbol{p}_2$, then T_3 is admissible as well, since it functions as a "second best" option. Convexity of the set of indeterminate probabilities, we note, is required in each of Gilboa and Schmeidler's (1989) version of Γ-*Maximin*, in Walley's (1990) version of *Maximality*, and in Levi's (1980) account of *E-admissibility*.

Aside: Example 2 relies on the fact that normal and extensive form decisions are generally *not* equivalent in decision theories with indeterminate probabilities. Example 2 is in the normal form, as are all the choice problems considered in this paper. In the extensive form of this decision problem, the decision maker has the opportunity to make a terminal choice between T_1 and T_2 first, or to take as a third option a sequential alternative: pay a fee of 0.01 utiles in order to learn the state of the weather before choosing between T_1 and T_2. Under decision rules for extensive form problems that we endorse, and which we believe also are endorsed by Levi, then it is *E-inadmissible* to postpone the immediate medical decision between T_1 and T_2 in order to pay an amount to acquire the irrelevant meteorological evidence. And this holds whether the indeterminate probability set is convex or not. Related results about independence with indeterminate probability are presented in Cozman and Walley (2005).

5 Concluding remarks

We have discussed *coherent choice functions*—where the admissible options in a decision problem are exactly those that maximize expected utility for some probability/utility pair in fixed set $\boldsymbol{S}$ of probability/utility pairs. All of the decision problems used here to characterize and axiomatize coherent choice functions are *normal form* decision problems. But, as indicated in Sect. 4, normal and extensive form decisions generally are not equivalent when probability (or utility) is indeterminate. One of our future projects is to study coherent choice for extensive form, i.e., sequential decision problems.

Also, as noted in Lemma 4, in parallel with our findings about coherent strict partial orders (1995), the axioms are sufficient for coherence using a set of probability/almost-state-independent utility pairs. Though we give sufficient conditions when a choice function satisfying our axioms is represented by a set of probability/state-independent utility pairs with a common utility, also we intend to study how to modify the axioms to avoid the use of almost-state-independent utilities.

Acknowledgements We thank Fabio Cozman, Isaac Levi (who sees Example 2 in a different light than do we), and Matthias Troffaes for very helpful discussions about topics central to this paper. Seidenfeld's research efforts were supported in part by State of Sao Paulo, Brazil, grant FAPESP #05/59541-4.

Appendix 1—Lemma 3

Lemma 2 *Suppose that $\boldsymbol{C}(\boldsymbol{O}^*) = \boldsymbol{O}^*$. Then $\boldsymbol{p}$ is a global Bayes model for the choice function $\boldsymbol{C}(\cdot)$ restricted to feasible sets in $\mathcal{H}_{\{0,1\}}$.*

Proof Let $\boldsymbol{p} = (p_1, \ldots, p_n)$ be a probability distribution on Ω with $\underline{\boldsymbol{p}}$ its smallest nonzero coordinate. We define $\boldsymbol{O}^*$ so that it is comprised by a set of acts that span all the elements of $\mathcal{H}_{\{0,1\}}$ with $\boldsymbol{p}$-expected utility $\underline{\boldsymbol{p}}$.

Partition the states in Ω in two sets:

$$\Omega_1^p = \{\omega_1, \ldots, \omega_k\} \text{ is the support of } \boldsymbol{p}$$

and,

$$\Omega_2^p = \{\omega_{k+1}, \ldots, \omega_n\} \text{ are those states null under } \boldsymbol{p}.$$

Clearly, $\Omega_2^p = \phi$ if and only if $\boldsymbol{p}$ has full support. We define $\boldsymbol{O}^*$ by two cases, depending whether $\Omega_2^p = \phi$ or not.

Case 1 $\Omega_2^p = \phi$ and $\boldsymbol{p}$ has full support. $\boldsymbol{O}^*$ is comprised by n-many acts, $\{a_j : j = 1, \ldots, n\}$. For each $\boldsymbol{j} = 1, \ldots, \boldsymbol{n}$, define the act a_j by

$$\begin{aligned} a_j(\omega_i) &= \frac{\underline{p}}{p_j}\mathbf{1} \oplus \left(1 - \frac{\underline{p}}{p_j}\right)\mathbf{0} \quad &\text{if} \quad i = j \\ a_j(\omega_i) &= 0 \quad &\text{if} \quad i \neq j. \end{aligned}$$

Case 2 $\Omega_2^p \neq \phi$. Then $\boldsymbol{O}^*$ is defined by $\boldsymbol{k}(\boldsymbol{n}+2-\boldsymbol{k})$-many acts which can be understood to be the product of acts defined on $\Omega_1^p \times \Omega_2^p$. With respect to Ω_1^p, $\boldsymbol{O}^*$ contains k-many acts that span horse lotteries defined on Ω_1^p that have $\boldsymbol{p}$-Expected utility $\underline{\boldsymbol{p}}$, similarly to Case 1. With respect to Ω_2^p, $\boldsymbol{O}^*$ contains $(\boldsymbol{n}+2-\boldsymbol{k})$-many acts that span all horse lotteries defined on Ω_2^p, including the two constants $\mathbf{0}$ and $\mathbf{1}$.

For each $\boldsymbol{j}=1,\ldots,\boldsymbol{k}$, and $\boldsymbol{m}=\boldsymbol{k}+1,\ldots,\boldsymbol{n}+2$ define the act a_j^m by

$$
\begin{array}{ll}
a_j^m(\omega_i) = \frac{\underline{p}}{p_j}\mathbf{1} \oplus \left(1-\frac{\underline{p}}{p_j}\right)\mathbf{0} & \text{if } \ i=j \\
a_j^m(\omega_i) = \mathbf{1} & \text{if } \ i=m \ \textbf{or} \ (m=n+2 \text{ and } i>k) \\
a_j^m(\omega_i) = \mathbf{0} & \text{otherwise}
\end{array}
$$

Note that $a_j^{n+1}(\omega_i) \neq \mathbf{0}$ if and only if $i=j$. In particular, it equals $\mathbf{0}$ on Ω_2^p. And note that $a_j^{n+2}(\omega_i) \neq \mathbf{0}$ if and only if, either $i=j$ or $i>k$. It equals $\mathbf{1}$ on Ω_2^p.

Let $\boldsymbol{O^*}$ be the choice problem formed by taking the convex hull of these options. That is, in Case 1 $\boldsymbol{O^*}=\boldsymbol{H}\left\{a_j : j=1,\ldots,n\right\}$, the convex hull of $\boldsymbol{n}$-many options. In Case 2, $\boldsymbol{O}^*=\boldsymbol{H}\left\{a_j^m : j=1,\ldots,\boldsymbol{k}; \boldsymbol{m}=\boldsymbol{k}+1,\ldots,\boldsymbol{n}+2\right\}$, the convex hull of $\boldsymbol{k}(\boldsymbol{n}+2-\boldsymbol{k})$-many options.

Let $\boldsymbol{a}_p$ denote the constant horse lottery that awards the identical von Neumann–Morgenstern lottery in each state, with

$$\boldsymbol{a_p} = \underline{\boldsymbol{p}}\mathbf{1} \oplus \left(1-\underline{\boldsymbol{p}}\right)\mathbf{0}.$$

Claim 1 $\boldsymbol{a_p} \in \boldsymbol{O}^*$.

Proof In Case 1, when $\boldsymbol{p}$ has full support, $\boldsymbol{p}_1a_1 \oplus \boldsymbol{p}_2a_2 \oplus \ldots \oplus \boldsymbol{p_n}\boldsymbol{a_n}$ is the horse lottery $\boldsymbol{a_p}$. In Case 2, when $\boldsymbol{p}$-null states exist, for each $\boldsymbol{j}=1,\ldots,\boldsymbol{k}$, define the horse lottery $b_{\boldsymbol{j}}=(1-\underline{\boldsymbol{p}})a_j^{n+1} \oplus \underline{\boldsymbol{p}}a_j^{n+2}$ with payoffs:

$$
\begin{array}{ll}
b_j(\omega_i) = \boldsymbol{a_p} & \text{if } i>k \\
b_j(\omega_i) = \frac{\underline{p}}{p_j}\mathbf{1} \oplus \left(1-\frac{\underline{p}}{p_j}\right)\mathbf{0} & \text{if } i=j \\
b_j(\omega_i) = \mathbf{0} & \text{if } i \neq j \text{ and } i \leq k.
\end{array}
$$

Then $\boldsymbol{p}_1b_1 \oplus \boldsymbol{p}_2b_2 \oplus \ldots \oplus \boldsymbol{p}_k b_{\mathbf{k}}$ is the horse lottery $\boldsymbol{a}_p$. □ − claim1.

Note that $(\boldsymbol{p},\boldsymbol{u})$ is a local Bayes model for each element of $\boldsymbol{O}^*$ as the $\boldsymbol{p}$-Expected utility for each element of $\boldsymbol{O}^*$ is the same value, namely $\underline{\boldsymbol{p}}$.

Claim 2 If $\underline{\boldsymbol{p}}<1$ then $(\boldsymbol{p},\boldsymbol{u})$ is the only local Bayes model for a_p

Proof Consider $\boldsymbol{q} \neq \boldsymbol{p}$. Regardless the distribution $\boldsymbol{q}$ on Ω, $\boldsymbol{a_p}$ has $\boldsymbol{q}$-Expected utility $\underline{\boldsymbol{p}}$. We argue by cases that when $\underline{\boldsymbol{p}}<1$, $\boldsymbol{q}$ is not a local model for $\boldsymbol{a_p}$ with respect to $\boldsymbol{O}^*$.

If $\boldsymbol{p}$ has full support ($\Omega_2^p = \phi$), the $\boldsymbol{q}$-Expected utility of $a_j = \boldsymbol{q}_j \frac{\underline{\boldsymbol{p}}}{\boldsymbol{p}_j} > \underline{\boldsymbol{p}}$. And if $j = m > k$, so that $\boldsymbol{p}_j = 0$ and $q(\Omega_2^p) > 0$, then the $\boldsymbol{q}$-Expected utility of $\boldsymbol{p}_1 a_1^{n+2} \oplus \boldsymbol{p}_2 a_2^{n+2} \oplus \ldots \oplus \boldsymbol{p}_k a_k^{n+2} = \boldsymbol{q}(\Omega_1)\underline{\boldsymbol{p}} + \boldsymbol{q}(\Omega_2) > \underline{\boldsymbol{p}}$.

Hence, $(\boldsymbol{q},\boldsymbol{u})$ is not a local Bayes model for $\boldsymbol{a_p}$. □−claim2.

Note also that for the case $\boldsymbol{p}_1 = \underline{\boldsymbol{p}} = 1$, $\boldsymbol{a_p} = \mathbf{1}$ and then $\boldsymbol{O}^* = H\left\{\mathbf{1}, a_1^2, \ldots a_1^{n+2}\right\}$. In which case, if $\boldsymbol{q} \neq \boldsymbol{p}$, $\boldsymbol{q}$ is not a local Bayes model for a_1^{n+1}, which has a $\boldsymbol{q}$-expected value of $\boldsymbol{q}_1 < 1$. Thus, we have

Proposition *$\boldsymbol{p}$ is the sole local Bayes model for all of $\boldsymbol{O}^*$.*

Claim 3 *$\boldsymbol{O}^*$ contains all the horse lotteries in $\mathcal{H}_{\{0,1\}}$ that have $\boldsymbol{p}$-expected utility equal to $\underline{\boldsymbol{p}}$.*

Proof Let o be such a horse lottery with $\boldsymbol{p}$-Expected utility $\underline{\boldsymbol{p}}$. Write $o(\omega_j) = \alpha_{\boldsymbol{j}}\mathbf{1} \oplus (1-\alpha_{\boldsymbol{j}})\mathbf{0}$, $\boldsymbol{j} = 1, \ldots, \boldsymbol{n}$.

Case 1 *($\boldsymbol{p}$ has full support.)*: For $\omega_i \in \Omega = \Omega_1^p$ we have that $\sum_i \boldsymbol{p}_i \alpha_i = \underline{\boldsymbol{p}}$ and $0 \le \alpha_i \le 1$. The set of α-vectors satisfying these two equations is closed and convex, with extreme points given by the acts $\{\boldsymbol{a}_j\}$. That is, if $\alpha^* = < \alpha *_1, \ldots, \alpha *_n >$ is an extreme point of this set of α-vectors, then $\alpha^* = \alpha_{\boldsymbol{j}}$ for some $1 \le \boldsymbol{j} \le \boldsymbol{n}$. Since a closed, convex set is identified by its extreme points, this establishes that $\boldsymbol{o} \in \boldsymbol{O}^*$.

Case 2 *(There are null states under $\boldsymbol{p}$.)*: The reasoning is similar to Case 1, noting that $\boldsymbol{O}^*$ spans all horse lotteries defined over Ω_2^p. □−claim3.

We complete the proof of Lemma 3, as follows. Let $\boldsymbol{O}$ be a choice set and let $\phi \neq \boldsymbol{O_p} \subseteq \boldsymbol{O}$ be those options for which $\boldsymbol{p}$ is a local Bayes model. So, each $a \in \boldsymbol{O_p}$ maximizes the $\boldsymbol{p}$-Expected utility of options in $\boldsymbol{O}$ at common value $\mathbf{r}$. There are two cases, depending upon whether $\mathbf{r} \ge \underline{\boldsymbol{p}}$ or $\mathbf{r} < \underline{\boldsymbol{p}}$.

In the former case, mix $\mathbf{0}$ into each act in $\boldsymbol{O}$ to form the choice set $\boldsymbol{O}' = \frac{\underline{p}}{\boldsymbol{r}}\boldsymbol{O} \oplus \left(1 - \frac{\underline{p}}{\boldsymbol{r}}\right)\mathbf{0}$, with the isomorphism between $\boldsymbol{O}$ and $\boldsymbol{O}'$ that associates each $o \in \boldsymbol{O}$ with $o' \in \boldsymbol{O}$, where $o' = \frac{\underline{p}}{\boldsymbol{r}} o \oplus \left(1 - \frac{\underline{\boldsymbol{p}}}{\boldsymbol{r}}\right)\mathbf{0}$.

In case $\mathbf{r} < \underline{\boldsymbol{p}}$ then mix $\mathbf{1}$ into each act in $\boldsymbol{O}$ to form the choice set $\boldsymbol{o}' = \frac{1-\underline{\boldsymbol{p}}}{1-\boldsymbol{r}}\boldsymbol{O} \oplus \left(\frac{1-\underline{\boldsymbol{p}}}{1-\boldsymbol{r}}\right)\mathbf{1}$, with the isomorphism between $\boldsymbol{O}$ and $\boldsymbol{o}'$ that associates each $o \in \boldsymbol{O}$ with $o' \in \boldsymbol{O}$, where $o' = \frac{1-\underline{\boldsymbol{p}}}{1-\boldsymbol{r}} o \oplus \left(\frac{1-\underline{\boldsymbol{p}}}{1-\boldsymbol{r}}\right)\mathbf{1}$.

The argument continues in parallel between the two cases. By the Axiom 2, $a \in \boldsymbol{C}(\boldsymbol{O})$ if and only if $a' \in \boldsymbol{C}(\boldsymbol{O}')$. Also evident is the fact that for each $a' \in \boldsymbol{O}'_{\boldsymbol{p}}$ the $\boldsymbol{p}$-Expected utility of a' equals $\underline{\boldsymbol{p}}$. Thus, by Claim 3, for each $a' \in \boldsymbol{O}'_{\boldsymbol{p}}$, $a' \in \boldsymbol{C}(\boldsymbol{O}^*)$.

Claim 4 Let $o' \in \boldsymbol{O}'$ and $o' \notin \boldsymbol{O}'_{\boldsymbol{p}}$. Then each local Bayes model $\boldsymbol{q}$ for o' with respect to $\boldsymbol{O}^* \cup \{o'\}$ is singular with respect to $\boldsymbol{p}$, i.e., $\Omega_1^q \cap \Omega_1^p = \phi$.

Proof Because $o' \notin \boldsymbol{O}'_{\boldsymbol{p}}$ then $\mathbf{E}_{\boldsymbol{p}}(o') < \underline{\boldsymbol{p}}$ and, trivially, $\boldsymbol{p}$ is not a local Bayes model for o'. Fix a distribution $\boldsymbol{q} \neq \boldsymbol{p}$ where $\Omega_1^q \cap \Omega_1^p \neq \phi$. We argue indirectly that $\boldsymbol{q}$ is not a local Bayes model for $\boldsymbol{o}'$ with respect to $\boldsymbol{O}^* \cup \{o'\}$.

First consider the case where $\Omega_1^q \subseteq \Omega_1^p$, that is where $\boldsymbol{q}$ is absolutely continuous with respect to $\boldsymbol{p}$. Within the $\boldsymbol{n}-1$ dimensional simplex of distributions on Ω, let $\boldsymbol{L_{pq}}$ be the line determined by the two points $\boldsymbol{p}$ and $\boldsymbol{q}$, having endpoints denoted $\boldsymbol{q}_*$ and $\boldsymbol{q}^*$. Identify these endpoints by placing $\boldsymbol{q}$ in the closed line segment $[\boldsymbol{q}^*, \boldsymbol{p}]$, and thus $\boldsymbol{p}$ lies in the closed line segment $[\boldsymbol{q}, \boldsymbol{q}*]$, from which we know that $\boldsymbol{p} \neq \boldsymbol{q}^*$, though it is possible that $\boldsymbol{q} = \boldsymbol{q}^*$.

Moreover, since $\Omega_2^q \supseteq \Omega_2^p$ we have that $\boldsymbol{p} \neq \boldsymbol{q}*$, since each endpoint of $\boldsymbol{L_{pq}}$ has some null-state not shared as a null state with any other point on that line. So, $\boldsymbol{p}$ is internal to the line $\boldsymbol{L_{pq}}$. And because $\boldsymbol{q}*$ is an endpoint of $\boldsymbol{L_{pq}}$, as just argued, $\Omega_2^{q^*} \cap \Omega_1^p \neq \phi$. Assume then that $\omega_k \in \Omega_2^{q^*} \cap \Omega_1^p$. Since $\boldsymbol{p}$ lies on the line $[\boldsymbol{q}*, \boldsymbol{q}^*]$, then $\omega_k \in \Omega_1^{q^*}$.

Consider the act a_k^{n+1} (or the act a_k if $\boldsymbol{p}$ has full support). Since $\mathbf{E}_{\boldsymbol{q}}(\boldsymbol{o}') \geq \mathbf{E}_{\boldsymbol{q}}(a_k^{n+1})$ and $\mathbf{E}_{\boldsymbol{p}}(\boldsymbol{o}') < \mathbf{E}_{\boldsymbol{p}}(a_k^{n+1}) = \underline{\boldsymbol{p}}$, there exists a unique distribution $\boldsymbol{r}_k$ situated on the line $\boldsymbol{L_{pq}}$ and between $\boldsymbol{p}$ and $\boldsymbol{q}$ (possibly with $\boldsymbol{r}_k = \boldsymbol{q}$), such that $\mathbf{E}_{\boldsymbol{r_k}}(\boldsymbol{o}') = \mathbf{E}_{\boldsymbol{r_k}}(a_k^{n+1})$. Because expected utility is linear in probability, for each distribution $\boldsymbol{t}$ in the half open interval $(\boldsymbol{r_k}, \boldsymbol{q}*]$, $\mathbf{E}_{\boldsymbol{t}}(\boldsymbol{o}') < \mathbf{E}_{\boldsymbol{t}}(a_k^{n+1})$. But $E_{q*}[a_k^{n+1}] = 0 > E_{q*}[o']$, which is a contradiction as no act has a negative expected value. This completes the argument when $\boldsymbol{q}$ is absolutely continuous with respect to $\boldsymbol{p}$.

Next, assume that $\Omega_1^q \cap \Omega_1^p \neq \phi$ and write $\boldsymbol{q}(\bullet) = \boldsymbol{q}(\bullet \mid \Omega_1^p)\boldsymbol{q}(\Omega_1^p) + \boldsymbol{q}(\bullet \mid \Omega_2^p)\boldsymbol{q}(\Omega_2^p)$, where $\boldsymbol{q}(\Omega_1^p) > 0$. So, $\boldsymbol{q}(\bullet|\Omega_1^p)$ is absolutely continuous with respect to $\boldsymbol{p}$.

$\mathbf{E}_{\boldsymbol{q}}(\cdot) = \mathbf{E}_q(\cdot|\Omega_1^p)\boldsymbol{q}(\Omega_1^p) + \mathbf{E}_q(\cdot|\Omega_2^p)\boldsymbol{q}(\Omega_2^p)$. Since $\boldsymbol{a}_k^{n+2} \in \boldsymbol{O}^*$ and $\mathbf{E}_q(\boldsymbol{o}') \geq \mathbf{E}_q(a_k^{n+2})$, it follows that $\mathbf{E}_q(\boldsymbol{o}'|\Omega_1^p) \geq \mathbf{E}_q(a_k^{n+2}|\Omega_1^p) = \mathbf{E}_q(a_k^{n+1}|\Omega_1^p)$. However, as $\boldsymbol{q}(\cdot|\Omega_1^p)$ is absolutely continuous with respect to $\boldsymbol{p}$, we have the same situation involving $\boldsymbol{q}(\cdot|\Omega_1^{\boldsymbol{p}})$ and $\boldsymbol{p}$ as when $\boldsymbol{q}$ is absolutely continuous with respect to $\boldsymbol{p}$, completing the proof. $\square$−claim4.

Next, we show that if there is a local Bayes model for o' with respect to $\boldsymbol{O}^* \cup \{o'\}$, then no element of $\boldsymbol{O}^*$ becomes inadmissible by adding option $\boldsymbol{o}'$.

Claim 5 Assume that $a \in \boldsymbol{C}(\boldsymbol{O}^*)$, $o' \in \boldsymbol{O}'$ but $o' \notin \boldsymbol{O}'_p$, and let o' have a local Bayes model $\boldsymbol{q}$ with respect to $\boldsymbol{O}^* \cup \{o'\}$. Then $a \in \boldsymbol{C}(\boldsymbol{O}^* \cup \{o'\})$.

Proof Assume the premise. In the light of Axiom 4 we are done proving Claim 5 if we identify an act $a^* \in \boldsymbol{O}^*$ such that a^* weakly dominates o'. This we do as follows.

By Claim 4, $\boldsymbol{q}$ is singular with respect to $\boldsymbol{p}$. Consider an act a_k^{n+2} for $\omega_k \in \Omega_1^p$.

Definition For $\boldsymbol{W} \subseteq \Omega$ and act o, define the "called-off" act $o|\boldsymbol{W}$ by:

$$o(\omega)|\boldsymbol{W} = o(\omega) \text{ for } \omega \in \boldsymbol{W}, \text{ and } o(\omega)|\boldsymbol{W} = \mathbf{0} \text{ otherwise.}$$

Write o' as an sum of three call-off acts $o' = o'|\Omega_1^q + o'|(\Omega_2^p \cap \Omega_2^q) + o'|\Omega_1^p$, and likewise for $a_k^{n+2} = a_k^{n+2}|\Omega_1^q + a_k^{n+2}|(\Omega_2^p \cap \Omega_2^q) + a_k^{n+2}|\Omega_1^p$. Because $a_k^{n+2}(\omega) = \mathbf{1}$ for $\omega \in \Omega_2^p$, then $a_k^{n+2}|\Omega_1^q$ weakly dominates $o'|\Omega_1^q$, and likewise $a_k^{n+2}|(\Omega_2^p \cap \Omega_2^q)$ weakly dominates $o'|(\Omega_2^p \cap \Omega_2^q)$. By Claim 4, $o'|\Omega_1^p$ fails to have a local Bayes model with respect to $\boldsymbol{O}^* \cup \{o'|\Omega_1^p\}$. So, by Lemma 2, there exists an option $b \in \boldsymbol{H}(\boldsymbol{O}^*)$ that uniformly dominates $o'|\Omega_1^p$. Let $a^* = a_k^{n+1}|\Omega_1^p + a_k^{n+1}|(\Omega_2^p \cap \Omega_2^q) + b|\Omega_1^p$. Then a^* weakly dominates o' and, as $\mathbf{E}_p[a^*] = \mathbf{E}_p[b|\Omega_1^q] = \underline{\boldsymbol{p}}$, we have $a^* \in \boldsymbol{O}*$. $\square$−claim5.

Assume that $a' \in \boldsymbol{C}(\boldsymbol{O}^*)$. Let $\boldsymbol{N}' = \{o' : o' \in \boldsymbol{O}'$ and $o' \notin \boldsymbol{O}'_p$ but $\boldsymbol{o}'$ has no local Bayes model with respect to $\boldsymbol{O}^* \cup \{o'\}\}$. Then by Lemma 2, $o' \in \boldsymbol{R}(\boldsymbol{O}^* \cup \boldsymbol{N}')$. By Axiom 1, as $a' \in \boldsymbol{C}(\boldsymbol{O}^*)$ then $a' \in \boldsymbol{C}(\boldsymbol{O}^* \cup \boldsymbol{N}')$. If $o' \in (\boldsymbol{O}' - \boldsymbol{N}')$ then, using Claim 5, $a' \in \boldsymbol{C}(\boldsymbol{O}^* \cup \boldsymbol{N}' \cup \{o'\})$.

By a simple induction on a well-ordering of $\boldsymbol{O}' - \boldsymbol{N}'$, then $a' \in \boldsymbol{C}(\boldsymbol{O}^* \cup \boldsymbol{N}' \cup (\boldsymbol{O}' - \boldsymbol{N}') = \boldsymbol{C}(\boldsymbol{O}^* \cup \boldsymbol{O}')$. By Axiom 1, if $a' \in \boldsymbol{O}'$ then $a' \in \boldsymbol{C}(\boldsymbol{O}')$. Finally, by Axiom 2, $a \in \boldsymbol{C}(\boldsymbol{O})$. □-Lemma 3

Appendix 2—Lemma 4

Lemma 4 *For each admissible option $o \in \boldsymbol{C}(\boldsymbol{O})$ at least one of its local Bayes models is a global Bayes model or else there is a set of probability/almost-state-independent utility pairs that serve as a global Bayes-model.*

Proof The next claim, which we use to establish Lemma 4, extends the idea of Axiom 4 to the strict partial order $\langle$.

Claim 6 Suppose that for option sets $\boldsymbol{A}$, $\boldsymbol{B}$ and $\boldsymbol{D}$, $\boldsymbol{B}\langle\boldsymbol{A}$ and $\boldsymbol{B} \cap \boldsymbol{C}(\boldsymbol{D}) \neq \phi$. Then $\boldsymbol{A} \cap \boldsymbol{C}(closure\{\boldsymbol{D} - (\boldsymbol{B} \cup \boldsymbol{A})\}) \neq \phi$.

Proof (indirect) Suppose that $\boldsymbol{A} \subseteq \boldsymbol{R}(closure\{\boldsymbol{D} - (\boldsymbol{B} \cup \boldsymbol{A})\})$. By Axiom 1 applied twice, $\boldsymbol{A} \subseteq \boldsymbol{R}(\boldsymbol{D} \cup \boldsymbol{A})$ and $\boldsymbol{A} \subseteq \boldsymbol{R}(\boldsymbol{D} \cup \boldsymbol{A} \cup \boldsymbol{B})$. Since $\boldsymbol{B}\langle\boldsymbol{A}$, likewise $\boldsymbol{B} \subseteq \boldsymbol{R}(\boldsymbol{D} \cup \boldsymbol{A} \cup \boldsymbol{B})$. Thus, $\boldsymbol{A} \cup \boldsymbol{B}\langle\boldsymbol{D}$. By transitivity, $\boldsymbol{B}\langle\boldsymbol{D}$ and therefore $\boldsymbol{B} \cap \boldsymbol{C}(\boldsymbol{D}) = \phi$. □-claim 6.

Given $o \in \boldsymbol{C}(\boldsymbol{O})$ and following the ideas we used in (1995, Definition 19), we introduce the notion of a ***target set*** $\boldsymbol{T}(o, \boldsymbol{O})$ of probability distributions for o with respect to choice problem $\boldsymbol{O}$. The target set for o is a subset of the local Bayes models for o which, we show, contains all of its global Bayes models. We demonstrate that whenever the target set includes a boundary point, that boundary point is a global Bayes model.

Given a probability distribution $\boldsymbol{p}$, recall the decision problem $\boldsymbol{O_p} = \left\{a^p, h_1^p, \ldots, h_n^p\right\}$ defined in Sect. 2. We state without proof that whenever $\boldsymbol{C}(\boldsymbol{O_p}) = \boldsymbol{O_p}$ then $\boldsymbol{C}(\boldsymbol{O}^*) = \boldsymbol{O}^*$ for $\boldsymbol{O}^*$ defined with respect to $\boldsymbol{p}$ as in Lemma 3, and so $\boldsymbol{p}$ is a global Bayes model.

Definition $\boldsymbol{T}(o, \boldsymbol{O}) = \{\boldsymbol{p} : \boldsymbol{p}$ is local Bayes model for o in choice problem $\boldsymbol{O}$ and $\left\{h_1^p, \ldots, h_n^p\right\} \subseteq C(\boldsymbol{O_p})\}$

Claim 7 $\boldsymbol{T}(o, \boldsymbol{O})$ is a non-empty, convex set.

Proof Without loss of generality, and to simplify the presentation, we establish the claim for a binary state space $\Omega = \{\omega_1, \omega_2\}$. Convexity is shown as follows. Note that for $\boldsymbol{p}$ defined by $\boldsymbol{p}(\omega_2) = 0$, $h_2^p \in \mathrm{C}(\boldsymbol{O_p})$, and for $\boldsymbol{p}$ defined by $\boldsymbol{p}(\omega_2) = 1$, $h_1^p \in \mathrm{C}(\boldsymbol{O_p})$. And by Claim 6, if $h_2^p \in \mathrm{C}(\boldsymbol{O_p})$, then for all distributions $\boldsymbol{q}$ with $\boldsymbol{q}(\omega_2) \leq \boldsymbol{p}(\omega_2)$ we have $h_2^q \in \mathrm{C}(\boldsymbol{O_q})$; and if $h_1^p \in \mathrm{C}(\boldsymbol{O_p})$, then for all distributions $\boldsymbol{q}$ with $\boldsymbol{q}(\omega_2) \geq \boldsymbol{p}(\omega_2)$ we have $h_1^q \in \mathrm{C}(\boldsymbol{O_p})$. In the general case, with more than 2 states, the same result follows by noting that $\boldsymbol{T}(o, \boldsymbol{O})$ is an intersection of half-planes.

We show that $\boldsymbol{T}(o, \boldsymbol{O})$ is non-empty by an indirect argument using the Archimedean axiom. So, assume that for each $\boldsymbol{p}$, $\boldsymbol{C}\left\{h_1^p, h_2^p\right\}$ is a unit set, and by the observation above, let $\boldsymbol{q}$ be the $\boldsymbol{lub}\left\{\boldsymbol{p}(\omega_2) : h_2^p \in \boldsymbol{C}\left\{h_1^p, h_2^p\right\}\right.$. There are two cases.

Case 1 $\left\{h_2^q\right\} = \boldsymbol{C}\left\{h_1^q, h_2^q\right\}$ So $\boldsymbol{q}(\omega_2) < 1$ and then $h_1^q \langle h_2^q$ and for all $\boldsymbol{p}(\omega_2) > q(\omega_2)$, $h_2^p \langle h_1^p$. But as $\boldsymbol{p}$ approaches $\boldsymbol{q}$, h_i^p converges to h_i^q for $i = 1, 2$. Then by Axiom 3, $h_1^q \langle h_1^q$, a contradiction.

Case 2 $\left\{h_1^q\right\} = \boldsymbol{C}\left\{h_1^q, h_2^q\right\}$. So $\boldsymbol{q}(\omega_2) > 0$ and then $h_2^q \langle h_1^q$ and for all $\boldsymbol{p}(\omega_2) < q(\omega_2)$, $h_1^p \langle h_2^p$. But as $\boldsymbol{p}$ approaches $\boldsymbol{q}$, h_i^p converges to h_i^q for $i = 1, 2$. Then by Axiom 3, we obtain the contradiction, $h_2^q \langle h_2^q$. □-claim 7.

To complete the proof of Lemma 4 there are two cases to consider.

Case 1 $\boldsymbol{T}(o, \boldsymbol{O})$ contains at least one of its boundary points. Suppose, e.g., that $\boldsymbol{q}$ is the $\boldsymbol{lub}\left\{p(\omega_2) : h_2^p \in \boldsymbol{C}\left\{h_1^p, h_2^p\right\}\right.$ and that $\boldsymbol{R}\left\{h_1^q, h_2^q\right\} = \phi$. Then for each $0 \leq x \leq 1$, $\boldsymbol{R}\left\{h_1^q, h_2^q, xh_1^q \oplus (1-x)h_2^q\right\} = \phi$, as the following reasoning establishes.

Assume that $\boldsymbol{q}(\omega_2) < 1$, or we are done. Then for all $\boldsymbol{p}(\omega_2) > \boldsymbol{q}(\omega_2)$, $h_2^p \langle h_1^p$ as before. For $0 < x \leq 1$, by Axiom 2, $h_2^p \langle xh_1^p \oplus (1-x)h_2^p$. As $\boldsymbol{p}$ approaches $\boldsymbol{q}$, by Axiom 3, then $xh_1^q \oplus (1-x)h_2^q \in C\left\{h_1^q, h_2^q, xh_1^q \oplus (1-x)h_2^q\right\}$, on pain of contradiction, otherwise, that $h_2^q \langle h_2^q$. The reasoning is similar if the target set $\boldsymbol{T}(o, \boldsymbol{O})$ is closed at the other end. Then, at each point $\boldsymbol{p}$ of closure for $\boldsymbol{T}(o, \boldsymbol{O})$, $\boldsymbol{R}(\boldsymbol{O_p}) = \phi$ and $\boldsymbol{p}$ is global Bayes model.

Case 2 If the target set is entirely open and there is no $\boldsymbol{p} \in \boldsymbol{T}(o, \boldsymbol{O})$ such that $\boldsymbol{R}(\boldsymbol{O_p}) = \phi$, we arrive at the parallel situation studied in Sect. IV.2 of our (1995). That situation is one where, first, a coherent choice function $\boldsymbol{C}$ is induced by a finite set $\boldsymbol{P}$ of linearly independent probabilities on Ω. The convex target sets for $\boldsymbol{C}$ include subsets of $\boldsymbol{P}$ as extreme points, i.e., $\boldsymbol{R}(\boldsymbol{O_p}) = \phi$ for each $\boldsymbol{p} \in \boldsymbol{P}$. Hence, $\boldsymbol{C}$ is represented by the set $\boldsymbol{P}$ of global Bayes models. Then, this choice function $\boldsymbol{C}$ is changed into another $\boldsymbol{C}^+$, which is formed by adding the strict preferences, associated with finitely many conditions of the form $\boldsymbol{T}(o, \boldsymbol{O}) \cap \boldsymbol{R}(\boldsymbol{O_p}) \neq \phi$. The results established in Section IV.2 of our (1995) show that then $\boldsymbol{C}^+$ satisfies the axioms. Also, those results show that in a neighborhood of the extreme points of the target sets for $\boldsymbol{C}$ there are sets of probability/almost-state-independent utility pairs that are local Bayes models for $\boldsymbol{C}$, and which then represent the choice function $\boldsymbol{C}^+$. These almost-state-independent utilities result by adding at least one new prize $\{\boldsymbol{r}\}$ to the two $\{\boldsymbol{0}, \boldsymbol{1}\}$ used to create the horse lotteries studied here. □-Lemma 4

Corollary *If for each choice problem* $\boldsymbol{O}$ *and* $o \in \boldsymbol{C}(\boldsymbol{O})$, *the target set* $\boldsymbol{T}(o, \boldsymbol{O})$ *includes at least one of its boundary points, then* $\boldsymbol{C}$ *is represented by a set of probability/state-independent utility pairs.*

References

Aizerman, M. A. (1985). New problems in general choice theory. *Social Choice and Welfare, 2*, 235–282.

Anscombe, F. J., & Aumann, R. J. (1963). A definition of subjective probability. *The Annals of Mathematical Statistics, 34*, 199–205.

Cozman, F. G., & Walley, P. (2005). Graphoid properties of epistemic irrelevance and independence. *Annals of Mathematics and Artificial Intelligence, 45*, 173–195.

de Finetti, B. (1974). *Theory of probability* (Vol. 1). London: Wiley.

Gilboa, I., & Schmeidler, D. (1989). Maxmin expected utility with non-unique prior. *Journal of Mathematical Economics, 18*, 141–153.

Kadane, J. B., Schervish, M. J., & Seidenfeld, T. (2004). A Rubinesque theory of decision. *IMS Lecture Notes Monograph, 45*, 1–11.

Kindler, J. (1983). A general solution concept for two person, zero sum games. *Journal of Optimization Theory and Applications, 40*, 105–119.

Levi, I. (1974). On indeterminate probabilities. *Journal of Philosophy, 71*, 391–418.

Levi, I. (1986). *Hard choices*. Cambridge: Cambridge University Press.

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

Savage, L. J. (1954). *The foundations of statistics*. New York: Wiley.

Schervish, M. J., Seidenfeld, T., Kadane, J. B., & Levi, I. (2003). Extensions of expected utility theory and some limitations of pairwise comparisons. In J.-M. Bernard, T. Seidenfeld, & M. Zaffalon (Eds.), *Proceedings of the third international symposium on imprecise probabilities and their applications* (pp. 496–510). Carleton Scientific.

Schervish, M. J., Seidenfeld, T., & Kadane, J. B. (2008). Proper scoring rules, dominated forecasts, and coherence. T.R. #865, Department of Statistics, Carnegie Mellon University.

Seidenfeld, T., Schervish, M. J., & Kadane, J. B. (1995). A representation of partially ordered preferences. *Annals of Statistics, 23*, 2168–2217.

Sen, A. (1977). Social choice theory: A re-examination. *Econometrica, 45*, 53–89.

Wald, A. (1950). *Statistical decision functions*. New York: Wiley.

Walley, P. (1990). *Statistical reasoning with imprecise probabilities*. London: Chapman and Hall.