

March 7, 2012

Characterization of Proper and Strictly Proper Scoring Rules for Quantiles

Mark J. Schervish, Joseph B. Kadane, and Teddy Seidenfeld
Carnegie Mellon University

Abstract

We give necessary and sufficient conditions for a scoring rule to be proper (or strictly proper) for a quantile if utility is linear, and the distribution is unrestricted. We also give results when the set of distributions is limited, for example, to distributions that have first moments.

Keywords: Elicitation.

1. Introduction. In subjective Bayesian practice, elicitation is a critical matter (Savage, 1971; O’Hagan, 1998; O’Hagan et al., 2006; Garthwaite, Kadane, and O’Hagan, 2005). Many papers (Kadane et al., 1980 and Garthwaite and Dickey, 1985, 1988) rely on elicitation of quantiles even if the parametric forms they target are expressed in terms of moments. There are two good reasons for this. First, moments are not easy to understand. Second, moments are sensitive to extreme outliers, and may even not exist. Quantiles, on the other hand, always exist and are more intuitive.

There are two principal methods of eliciting quantiles, the direct method and by use of scoring rules. In the direct method, the expert or other person being elicited is asked, for example, “what value would leave you indifferent between betting that the outcome will be greater than the value you name and less than the value you name?”, to obtain a median (see Dey and Liu, 2007). Other values are then obtained by bisection.

Scoring rules, by contrast, give the expert an explicit penalty that is a function of the elicited quantile and the outcome. (Some authors, such as Gneiting and Raftery, 2007, define scoring rules to be gains to the expert rather than penalties, in which case a minus sign will convert results from one definition to the other.) The implicit assumption is that the expert’s utility is linear in the score. For example, absolute error is often thought of as a scoring rule for the median. A scoring rule is called *proper* if it is minimized at the desired quantile. It is *strictly proper* if it is minimized only at the desired quantile. Gneiting and Raftery (2007), following Cervera and Munoz (1996), propose a specific class of scoring rules that they prove to be proper for a certain restricted class of distributions. Additionally they write “We do not know whether [this class] provides the general form of proper scoring rule for quantiles.”

To address this issue, there are three problems to overcome. The first is that the condition that a scoring rule be proper is too weak to be useful. For example, a constant function is a proper scoring rule (because the desired quantile minimizes it), but so does every other

possible elicited value. Hence the main focus has to be on strictly proper scoring rules. Second, quantiles are not necessarily unique. For example, the pdf

$$f(x) = \begin{cases} 1/2 & 0 < x < 1 \\ 1/2 & 2 < x < 3 \\ 0 & \text{otherwise} \end{cases}$$

has the whole interval $[1,2]$ as medians. Finally, to be valid no matter what the underlying distribution is, one has to be careful about infinite expected scores. For example, absolute error under a Cauchy distribution has infinite expected score for all choices of median, making it not strictly proper.

In response to these challenges, we focus on strictly proper scoring rules (although we find a characterization of proper scoring rules as well). We define quantiles in a way that yield a closed interval in general (and specializes to a point in the case of an absolutely continuous distribution with positive density at the quantile). We extend the notion of proper scoring rule to mean that every point in this closed interval minimizes expected score, and the notion of strictly proper scoring rule to mean that only these points do so. Finally, our main result is valid for the class of all distributions, although we also give results for classes of distributions having certain finite moments.

2. Characterization of proper scoring rules. In this section, we give a characterization of (strictly) proper scoring rules for a single quantile in Theorem 1. Throughout this discussion, $0 < \alpha < 1$, P is a probability distribution over a set \mathcal{X} of real numbers, and X is a random variable whose distribution is P . Let $F(x) = P(X \leq x)$ stand for the c.d.f. of X , and let $G(x) = P(X \geq x)$ be the symmetrically defined function for the upper tail. It is useful to note that, for all $a < b$,

$$(1) \quad P(a < X < b) = 1 - G(b) - F(a).$$

We use the functions F and G to define quantiles.

DEFINITION 1. Let $0 < \alpha < 1$. Every real number q such that

$$(2) \quad F(q) \geq \alpha, \text{ and } G(q) \geq 1 - \alpha,$$

is called an α quantile of P and an α quantile of X .

There are other definitions of quantiles, some of which provide for a unique α quantile for each α , others of which allow for multiple quantiles, as does Definition 1. We choose the definition above because it is symmetric with respect to the two tails of the distribution. That is, q is an α quantile of X if and only if $-q$ is a $1 - \alpha$ quantile of $-X$. Some results are easier to state or understand in terms of random variables while others are easier to state or understand in terms of distributions. This is why Definition 1 defines quantiles both for random variables and for distributions. If h is an extended-real-valued function defined on \mathbb{R} , we use the symbols $P[h(X)]$ and $\int h(x)dP(x)$ interchangeably depending on whether or not an explicit random variable X is understood from context.

The following lemma summarizes some useful facts about quantiles.

LEMMA 1. Let X be a random variable with distribution P . The set of α quantiles of X is a nonempty closed interval $[q_\alpha, q^\alpha]$. Also,

- If $q^\alpha > q_\alpha$, then $G(q) = 1 - \alpha$ for all $q \in (q_\alpha, q^\alpha]$ and $F(q) = \alpha$ for all $q \in [q_\alpha, q^\alpha)$.
- If $F(q) = \alpha$ and $t > q^\alpha$, then $P[q < X < t] > 0$.
- If $G(q) = 1 - \alpha$ and $t < q_\alpha$, then $P[t < X < q] > 0$.

PROOF. Define $A = \{x : F(x) \geq \alpha\}$ and $B = \{x : G(x) \geq 1 - \alpha\}$. Define $q_\alpha = \inf A$ and $q^\alpha = \sup B$. Since $F(x) \geq F(q_\alpha)$, for all $x > q_\alpha$, it follows that $F(q_\alpha) \geq \alpha$. Similarly, $G(q^\alpha) \geq 1 - \alpha$. Note that both A and B are semi-infinite intervals, and both q_α and q^α are finite. Every $x < q_\alpha$ is in B and every $x > q^\alpha$ is in A , so that $q_\alpha \leq q^\alpha$. The set of quantiles is $A \cap B$, which is a closed and bounded interval.

For the first bullet, if $F(q_\alpha) > \alpha$, no number greater than q_α is in B , hence a necessary condition for $q^\alpha > q_\alpha$ is $F(q_\alpha) = \alpha$. Similarly, a necessary condition for $q^\alpha > q_\alpha$ is $G(q^\alpha) = 1 - \alpha$. For every $t < u$ in the open interval (q_α, q^α) , $F(t) = F(q_\alpha)$ and $G(u) = G(q^\alpha)$, hence the first bullet holds.

For the second bullet, it is clear that $G(t) < 1 - \alpha$ for all $t > q^\alpha$. Then (1) implies

$$P(q < X < t) = 1 - G(t) - F(q) > 1 - (1 - \alpha) - \alpha = 0.$$

The third bullet is proven in similar fashion, since $F(t) < \alpha$ for all $t < q_\alpha$. \square

Let $g_\alpha : \mathcal{X} \times \mathcal{R} \rightarrow \mathbb{R}$ be a function, where \mathcal{R} is the set of allowed values for the quantile of interest, e.g., the convex hull of \mathcal{X} , or even its closure. Suppose that we want to use $g_\alpha(x, q)$ as the penalty to an elicitee for giving q as the α quantile of P when $X = x$ is observed. If the goal is to elicit an α quantile, we would like the penalty to incentivize the elicitee to provide an α quantile for the answer.

DEFINITION 2. Let \mathcal{P}_0 be a collection of distributions over \mathbb{R} . We say that g_α is a *proper scoring rule* for the α quantiles of \mathcal{P}_0 if the following holds. For every $P \in \mathcal{P}_0$ and for each α quantile q^* of P , $P[g_\alpha(X, q)]$ exists and is minimized by $q = q^*$. We say that g_α is *strictly proper* if the α quantiles of P are the only values of q that minimize $P[g_\alpha(X, q)]$.

EXAMPLE 1. Let \mathcal{P}_0 consist of all distributions for random variables with finite mean. It is well known that $g_{1/2}(x, q) = |x - q|$ is a strictly proper scoring rule for the medians of \mathcal{P}_0 .

A necessary condition for g_α to be a (strictly) proper scoring rule for the α quantiles of a set \mathcal{P}_0 is that it be (strictly) proper for the α quantiles of each subset of \mathcal{P}_0 , in particular, those distributions in \mathcal{P}_0 that are supported on at most two values, if any. Another necessary condition for g_α to be strictly proper is that, for every $P \in \mathcal{P}_0$ there exists q such that $P[g_\alpha(X, q)] < \infty$.

LEMMA 2. Let \mathcal{P}_0 consist of all distributions supported on the set $\mathcal{X} = \{a, b\}$. A scoring rule $g_\alpha : \mathcal{X} \times [a, b] \rightarrow \mathbb{R}$ is proper for the α quantiles of \mathcal{P}_0 if and only if

- there is a number $d(a, b, \alpha)$ such that

$$(3) \quad g_\alpha(b, q) = d(a, b, \alpha) - \frac{\alpha}{1 - \alpha} g_\alpha(a, q),$$

for all $q \in [a, b]$, and

- $g_\alpha(a, a) \leq g_\alpha(a, b)$.

The scoring rule is strictly proper for the α quantiles of \mathcal{P}_0 if and only if, in addition to the above conditions, $g_\alpha(a, a) < g_\alpha(a, b)$.

PROOF. Let P_α be the distribution such that $P_\alpha(X = a) = \alpha$. Every number in the interval $[a, b]$ is an α quantile of X . A necessary condition for g_α to be proper is that $P_\alpha[g_\alpha(X, q)]$ be constant as a function of q . Trivially,

$$(4) \quad P_\alpha[g_\alpha(X, q)] = \alpha g_\alpha(a, q) + (1 - \alpha) g_\alpha(b, q).$$

This is constant in q if and only if there is a number $d(a, b, \alpha)$ such that (3) holds for all $q \in [a, b]$. Another way to understand (3) is that, aside from a shift of level, $\alpha g_\alpha(a, \cdot)$ and $(1 - \alpha) g_\alpha(b, \cdot)$ are curves with slopes that are negatives of each other for all $q \in [a, b]$. Another necessary condition for g_α to be proper comes from consideration of P_p , the distribution such that $P_p(X = a) = p$.

$$\begin{aligned} P_p[g_\alpha(X, q)] &= p g_\alpha(a, q) + (1 - p) g_\alpha(b, q) \\ &= p g_\alpha(a, q) + (1 - p) \left[d(a, b, \alpha) - \frac{\alpha}{1 - \alpha} g_\alpha(a, q) \right] \\ &= g_\alpha(a, q) \left[p - \frac{(1 - p)\alpha}{1 - \alpha} \right] + (1 - p) d(a, b, \alpha). \end{aligned}$$

The minimum of this expression, as a function of q , occurs at the minimum or maximum for $g_\alpha(a, \cdot)$ depending on whether the coefficient $p - (1 - p)\alpha/(1 - \alpha)$ is positive or negative respectively. The coefficient is positive if and only if $p > \alpha$, and it is negative if and only if $p < \alpha$. Hence, a necessary condition for g_α to be proper is $g_\alpha(a, a) \leq g_\alpha(a, b)$ with strict inequality being necessary for strict propriety. The steps used to prove that the conditions are necessary also show that the conditions are sufficient. \square

Now, let \mathcal{P}_0 contain all distributions supported on two-points (and possibly other distributions). A necessary condition for g_α to be (strictly) proper for the α quantiles of \mathcal{P}_0 is that the conditions stated in Lemma 2 hold simultaneously for all $a < b$. If, for each $P \in \mathcal{P}_0$, $P[g_\alpha(X, q)]$ has the same value (whether finite or infinite) for all q , then g_α is trivially a proper scoring rule. In particular, a function that does not depend at all on q is a proper scoring rule.

Of course, functions of the form $g_\alpha(x, q) = f(x)$ (including constant functions) will be proper for every set \mathcal{P}_0 of distributions P such that $P[f(X)]$ is defined. Such scoring rules are not strictly proper, and they are about as uninteresting as one could imagine. Nevertheless, such rules will be special cases of the following theorem (with $s(\cdot)$ being a constant function.)

THEOREM 1. *Let \mathcal{P}_0 be a collection of probability distributions on a subset \mathcal{X} of \mathbb{R} such that, for every $a, b \in \mathcal{X}$, \mathcal{P}_0 contains every distribution concentrated on $\{a, b\}$. Let g_α be a real-valued function defined on $\mathcal{X} \times \mathcal{R}$, where \mathcal{R} is the convex hull of \mathcal{X} . Then g_α is a (strictly) proper scoring rule for the α quantiles of \mathcal{P}_0 if and only if for each $P \in \mathcal{P}_0$ $P[g_\alpha(X, q)]$ exists (is finite), and there exists a (strictly) increasing function s such that*

$$(5) \quad g_\alpha(x, q) - g_\alpha(x, x) = \begin{cases} \alpha[s(x) - s(q)] & \text{if } x > q, \\ (1 - \alpha)[s(q) - s(x)] & \text{if } x < q. \end{cases}$$

PROOF. First, we prove necessity by assuming that g_α is (strictly) proper for the set of all two-point distributions. This allows us to assume that g_α satisfies the two bullets in Lemma 2. Define $g^*(x, q) = g_\alpha(x, q) - g_\alpha(x, x)$. The second bullet of Lemma 2, namely that $g_\alpha(a, a) \leq g_\alpha(a, b)$, when applied to all $a < b$ implies, among other things, that $g_\alpha(a, q)$ and $g^*(a, q)$ are monotone increasing in q for $q > a$ and monotone decreasing in q for $q < a$ (with strict monotonicity in the strictly proper case). From this it follows that $g^*(x, q) \geq 0$ for all x, q .

Rewrite (3) as

$$(6) \quad d(a, b, \alpha) = g_\alpha(b, q) + \frac{\alpha}{1 - \alpha} g_\alpha(a, q),$$

for all $a \leq q \leq b$. Substitute $q = a$ and $q = b$ on the right side of (6) and set the two results equal to each other to obtain

$$(7) \quad g_\alpha(b, a) + \frac{\alpha}{1 - \alpha} g_\alpha(a, a) = g_\alpha(b, b) + \frac{\alpha}{1 - \alpha} g_\alpha(a, b),$$

which can be rewritten as

$$(8) \quad (1 - \alpha)g_\alpha^*(b, a) = \alpha g_\alpha^*(a, b).$$

In other words, specifying $g_\alpha^*(x, q)$ for $x > q$ determines its values for $x < q$ by (8).

Let $r(q) = g_\alpha^*(0, q)$, which we have already shown is (strictly) increasing for $q > 0$ and (strictly) decreasing for $q < 0$. It is more convenient to work with a monotone function such as

$$(9) \quad s(q) = \begin{cases} \frac{r(q)}{1 - \alpha} & \text{if } q \geq 0, \\ -\frac{r(q)}{\alpha} & \text{if } q < 0, \end{cases}$$

which is (strictly) increasing. Next, set the right sides of (6) and (7) equal to each other:

$$g_\alpha(b, q) + \frac{\alpha}{1 - \alpha} g_\alpha(a, q) = g_\alpha(b, b) + \frac{\alpha}{1 - \alpha} g_\alpha(a, b),$$

which rearranges to become

$$(10) \quad g_\alpha^*(b, q) = \frac{\alpha}{1 - \alpha} [g_\alpha^*(a, b) - g_\alpha^*(a, q)],$$

for $a \leq q \leq b$. Hence, for $0 \leq q \leq x$, we have from (10) (with $a = 0$ and $b = x$),

$$g_\alpha^*(x, q) = \frac{\alpha}{1 - \alpha} [r(x) - r(q)] = \alpha[s(x) - s(q)].$$

For $q \leq 0 \leq x$, we have from (10),

$$g_\alpha^*(x, 0) = \frac{\alpha}{1 - \alpha} [g_\alpha^*(q, x) - g_\alpha^*(q, 0)],$$

Rearranging terms and using (8) and (9), we get

$$g_\alpha^*(x, q) = \alpha[s(x) - s(q)].$$

For $q \leq x \leq 0$, we have from (10),

$$g_\alpha^*(0, x) = \frac{\alpha}{1 - \alpha} [g_\alpha^*(q, 0) - g_\alpha^*(q, x)],$$

from which it follows that

$$(11) \quad g_\alpha^*(x, q) = \alpha[s(x) - s(q)].$$

Hence, we see that (11) holds whenever $x \geq q$. When $x < q$, we apply (8) to get

$$(12) \quad g_\alpha^*(x, q) = (1 - \alpha)[s(q) - s(x)].$$

The condition that $P[g_\alpha(X, q)]$ exists is part of the definition of proper scoring rule. If g_α is strictly proper, it is necessary that the minimum of $P[g_\alpha(X, q)]$ occur only at α quantiles, hence it is neither possible for $P[g_\alpha(X, q)] = \infty$ for all q nor possible for $P[g_\alpha(X, q)] = -\infty$ for all q . For each real t , define $h_t^*(x, q) = g_\alpha(x, q) - g_\alpha(x, t)$. Then

$$(13) \quad h_t^*(x, q) = \begin{cases} (1 - \alpha)[s(q) - s(t)] & \text{if } x \leq \min\{t, q\}, \\ s(x) - \alpha s(q) - (1 - \alpha)s(t) & \text{if } q < x < t, \\ (1 - \alpha)s(q) + \alpha s(t) - s(x) & \text{if } t < x < q, \\ \alpha[s(t) - s(q)] & \text{if } x \geq \max\{t, q\}. \end{cases}$$

Note that $h_t^*(x, q)$, which equals the difference between the scoring rule at two possible quantiles q and t , is bounded as a function of x for each q and t . Hence, $P[g_\alpha(X, q)]$ is infinite for some q if and only if it is infinite (with the same sign) for all q . So, if g_α is strictly proper, then $P[g_\alpha(X, q)]$ is finite for all q . This completes the necessity part of the proof.

Next, we prove sufficiency by showing that all g_α with the specified form are (strictly) proper. Let X be a random variable, with probability distribution $P \in \mathcal{P}_0$. If $P[g_\alpha(X, t)]$ is infinite for some t , then $P[g_\alpha(X, q)]$ is infinite (the same sign) for all q , and g_α might be proper, but not strictly proper. For the remainder of the proof, assume that $P[g_\alpha(X, t)]$ is finite for all t . It now suffices to show that $P[g_\alpha(X, q)]$ is (uniquely) minimized when q is an α quantile of P .

Because $h_t^*(x, q) = g_\alpha(x, q) - g_\alpha(x, t)$, $P[g_\alpha(X, q)]$ is (uniquely) minimized at α quantiles of P if and only if the following are both true:

1. If q and t are both quantiles, then $P[h_t^*(X, q)] = 0$.

2. If q is a quantile and t is not a quantile, then $P[h_t^*(X, q)] \leq 0$, with strict inequality for unique minimization.

Let q be an α quantile of X , and let $t < q$. Then, we can write

$$\begin{aligned}
P[h_t^*(X, q)] &= [s(q) - s(t)][(1 - \alpha)F(t) - \alpha G(q)] + [(1 - \alpha)s(q) + \alpha s(t)]P[t < X < q] \\
&\quad - \int_{(t, q)} s(x)dP(x) \\
(14) \quad &= s(t)[\alpha - F(t)] + s(q)[1 - \alpha - G(q)] - \int_{(t, q)} s(x)dP(x).
\end{aligned}$$

First, assume that t is also an α quantile of P . Because $t < q$, $q > q_\alpha$ and $t < q^\alpha$ so that $F(t) = \alpha$ and $G(q) = 1 - \alpha$ according to Lemma 1, and $P[t < X < q] = 0$. This establishes condition 1.

Next, let q be an α quantile of X , and let $t < q$ not be an α quantile. In this case, $G(q) \geq 1 - \alpha$, and

$$(15) \quad \int_{(t, q)} s(x)dP(x) \geq s(t)[1 - F(t) - G(q)].$$

There are two cases. (i) If $G(q) = 1 - \alpha$, then Lemma 1 says that $P(t < X < q) > 0$, which implies that the inequality in (15) is strict. It follows that (14) is strictly less than $[1 - \alpha - G(q)][s(q) - s(t)] = 0$. (ii) If $G(q) > 1 - \alpha$, then (14) is at most $[1 - \alpha - G(q)][s(q) - s(t)] \leq 0$ with strict inequality if s is strictly increasing. This establishes condition 2.

Next, let q be an α quantile of X , and let $t > q$. Then, we can write

$$\begin{aligned}
P[h_t^*(X, q)] &= [s(q) - s(t)][(1 - \alpha)F(q) - \alpha G(t)] - [(1 - \alpha)s(t) + \alpha s(q)]P[q < X < t] \\
&\quad + \int_{(q, t)} s(x)dP(x) \\
(16) \quad &= -s(q)[\alpha - F(q)] - s(t)[1 - \alpha - G(t)] + \int_{(q, t)} s(x)dP(x).
\end{aligned}$$

First, assume that t is also an α quantile of P . Because $t > q$, $t > q_\alpha$ and $q < q^\alpha$ so that $F(q) = \alpha$ and $G(t) = 1 - \alpha$ according to Lemma 1, and $P[q < X < t] = 0$. This establishes condition 1.

Finally, let q be an α quantile of X , and let $t > q$ not be an α quantile. In this case, $F(q) \geq \alpha$, and

$$(17) \quad \int_{(q, t)} s(x)dP(x) \leq s(t)[1 - F(q) - G(t)].$$

There are two cases. (i) If $F(q) = \alpha$, then Lemma 1 says that $P(q < X < t) > 0$, which implies that the inequality in (17) is strict. It follows that (16) is strictly greater than $[\alpha - F(q)][s(t) - s(q)] = 0$. (ii) If $F(q) > \alpha$, then (16) is at most $[\alpha - F(q)][s(t) - s(q)] \leq 0$ with strict inequality if s is strictly increasing. This establishes condition 2. \square

EXAMPLE 2. The scoring rule $g_{1/2}(x, q) = |x - q|$ in Example 1 corresponds to $s(x) = 2x$ in Theorem 1 with $g_{1/2}(x, x) \equiv 0$. If desired, one could extend this example to all distributions by setting $g_{1/2}(x, x) = -|x|$. The scoring rule would no longer be $|x - q|$, but $g_{1/2}(x, q) - g_{1/2}(x, t)$ would be the same as it would be in Example 1 for all x, q, t .

The negatives of the scoring rules in Theorem 6 of Gneiting and Raftery (2007) have the form given in Theorem 1.

3. The class of probability distributions. Theorem 1 is stated with minimal assumptions on the the class \mathcal{P}_0 of probability distributions for which the scoring rules are intended to be (strictly) proper. The only assumptions are that all distributions supported on two values are included and that the proposed scoring rule is real-valued. The definition of (strictly) proper scoring rule imposes conditions on the mean of the scoring rule. If we make different assumptions about \mathcal{P}_0 , other scoring rules can be (strictly) proper.

EXAMPLE 3. Let \mathcal{P}_0 be the collection of symmetric unimodal distributions with finite variance. Then each $P \in \mathcal{P}_0$ has a unique median that equals the mean. Hence, $g_{1/2}(x, q) = (x - q)^2$ is strictly proper for the medians of \mathcal{P}_0 . Note that distributions concentrated on two distinct values can't be both symmetric and unimodal, so Theorem 1 does not apply to this example.

The conditions in Theorem 1 on the mean of $g_\alpha(X, q)$ are necessary, but some light can be shed on them. Take, for example, the condition that $P[g_\alpha(X, q)]$ be finite for all $P \in \mathcal{P}_0$, which is necessary for g_α to be strictly proper. In Example 2, we noticed that $|x - q|$ is strictly proper for the class of all distributions that have finite mean for X . What is the entire collection of strictly proper scoring rules for this class of distributions? The following result helps.

LEMMA 3. *Let f and h be real-valued functions of a real argument. Let \mathcal{P}_0 be the set of all probability distributions on \mathbb{R} that give finite mean to $f(X)$. Then $P[h(X)]$ is finite for all $P \in \mathcal{P}_0$ if and only if*

$$(18) \quad \sup_x \frac{|h(x)|}{1 + |f(x)|} < \infty.$$

PROOF. For the “if” direction, assume that (18) is true. Let M be the supremum in (18). Then

$$P[|h(X)|] \leq MP[1 + |f(X)|] = M + MP[|f(X)|],$$

which is finite for every $P \in \mathcal{P}_0$. For the “only if” direction, assume that (18) is false. We will find a $P \in \mathcal{P}_0$ such that $P[|h(X)|] = \infty$. Let x_1 be such that $|h(x_1)|/[1 + |f(x_1)|] \geq 1$. For each $n > 1$, let $x_n \notin \{x_1, \dots, x_{n-1}\}$ be such that $|h(x_n)|/[1 + |f(x_n)|] \geq n$. Let

$$P(\{x_n\}) = \frac{c}{[1 + |f(x_n)|]n^2},$$

where c is chosen to make $\sum_{n=1}^{\infty} P(\{x_n\}) = 1$. Then

$$P[|f(X)|] = \sum_{n=1}^{\infty} \frac{c|f(x_n)|}{[1 + |f(x_n)|]n^2} < \infty,$$

so that $P \in \mathcal{P}_0$. Also,

$$P[|h(X)|] = \sum_{n=1}^{\infty} \frac{c|h(x_n)|}{[1 + |f(x_n)|]n^2} \geq \sum_{n=1}^{\infty} \frac{cn[1 + |f(x_n)|]}{[1 + |f(x_n)|]n^2} = \infty. \quad \square$$

Lemma 3 lets us identify all strictly proper scoring rules for classes of distributions defined by certain means being finite.

THEOREM 2. *Let \mathcal{F} be a collection of real-valued functions defined on \mathbb{R} . Let \mathcal{P}_0 be the set of all probability distributions on \mathbb{R} that give finite mean to $f(X)$ for every $f \in \mathcal{F}$. Define \mathcal{H} to be the set of all functions h that satisfy (18) for all $f \in \mathcal{F}$. Then g_α is a strictly proper scoring rule for the α quantiles of \mathcal{P}_0 if and only if there exists $h \in \mathcal{H}$ and a strictly increasing function s such that, for all real t , $g_\alpha(x, q) = h(x) + h_t^*(x, q)$, where h_t^* is defined in (13).*

PROOF. First note that all distributions supported on two points are in \mathcal{P}_0 , no matter what \mathcal{F} is. For the “if” direction, we already showed in the proof of Theorem 1 that every h_t^* of the form (13) is strictly proper for the class of all probability measures. Adding a function that has finite mean for all $P \in \mathcal{P}_0$ produces another strictly proper scoring rule for \mathcal{P}_0 . For the “only if” direction, assume that g_α is strictly proper. It follows that $P[g_\alpha(X, t)]$ is finite for every $P \in \mathcal{P}_0$ and every real t , hence $h(x) = g_\alpha(x, t) \in \mathcal{H}$. In the proof of Theorem 1, we showed that there is a strictly increasing s such that, for every real t , $g_\alpha(x, q) - g_\alpha(x, t) = h_t^*(x, q)$ from (13). Hence $g_\alpha(x, q) = h(x) + h_t^*(x, q)$. \square

4. Conclusion. This paper characterizes proper and strictly proper scoring rules for a quantile if utility is linear in score. Kiefer (2010) following Karni (2009) finds a proper scoring rule under more general risk-averse utility. Characterizing proper and strictly proper scoring rules under those conditions remains open.

It is obvious that the sum of (strictly) proper scoring rules for several quantities is (strictly) proper. But are there others? The characterization of such rules is also currently unsolved.

REFERENCES

- Cervera, J. and Munoz, J. (1996), “Proper scoring rules for fractiles,” in *Bayesian Statistics 5*, eds. J. Bernardo, J. Berger, A. Dawid, and A. Smith, U.K.: Oxford University Press, pp. 513–519.

- Dey, D. and Liu, J. (2007), “A quantitative study of quantile based direct prior elicitation from expert opinion,” *Bayesian Analysis*, 2, 137–166.
- Garthwaite, P. and Dickey, J. (1985), “Double- and single-bisection methods for subjective probability assessment in a location-scale family,” *Journal of Econometrics*, 29, 149–163.
- (1988), “Quantifying expert opinion in linear regression problems,” *Journal of the Royal Statistical Society*, 50(Series B), 462–474.
- Garthwaite, P., Kadane, J., and O’Hagan, A. (2005), “Statistical Methods for eliciting probability distributions,” *Journal of the American Statistical Association*, 100, 680–701.
- Gneiting, T. and Raftery, A. (2007), “Strictly Proper Scoring Rules, Prediction and Estimation,” *Journal of the American Statistical Association*, 102, 359–378.
- Kadane, J., Dickey, J., Winkler, R., Smith, W., and Peters, S. (1980), “Interactive elicitation from the normal linear model,” *Journal of the American Statistical Association*, 75, 845–854.
- Karni, E. (2009), “A mechanism for eliciting probabilities,” *Econometrica*, 77, 603–606.
- Kiefer, N. (2010), “Incentive-compatible elicitation of quantiles,” unpublished manuscript.
- O’Hagan, A. (1998), “Eliciting expert beliefs in substantial practical applications,” *The Statistician*, 47, 21–35.
- O’Hagan, A., Buck, C., Daneshkhah, A., Eiser, J., Garthwaite, P., Jenkinson, D., Oakley, J., and Rakow, T. (2006), *Uncertain Judgement: Eliciting Expert Probabilities*, Chichester: J. Wiley and Sons.
- Savage, L. (1971), “Elicitation of Personal Probabilities and Expectations,” *Journal of the American Statistical Association*, 66, 783–801.