

Scoring Rules, Entropy, and Imprecise Probabilities

Robert Nau, Victor Richmond Jose, and Robert Winkler

Fuqua School of Business, Duke University

robert.nau@duke.edu

Abstract

Suppose that a risk-averse expected utility maximizer with a precise probability distribution \mathbf{p} bets optimally against a risk neutral opponent (or equivalently invests in an incomplete market for contingent claims) whose beliefs (or prices) are described by a convex set \mathcal{Q} of probability distributions. This utility-maximization problem is the dual of the problem of finding the distribution \mathbf{q} in \mathcal{Q} that minimizes a generalized divergence (relative entropy) with respect to \mathbf{p} . A special case is that of logarithmic utility, in which the corresponding divergence is the Kullback-Leibler divergence, but we present a closed-form solution for the entire family of linear-risk-tolerance (a.k.a. HARA) utility functions and show that this corresponds to a particular parametric family of generalized divergences, which is derived from an entropy measure originally proposed by Arimoto and which is also related to a generalization of pseudospherical scoring rule originally proposed by I.J. Good. A variant of this decision problem, in which the decision maker has quasilinear utility for consumption over two periods, leads to the family of power divergences, which is related to a generalization of the power family of scoring rules.

Keywords. entropy, divergence, scoring rules, portfolio optimization, incomplete markets

1 Introduction

There are many applications in which it is of interest to measure the difference between a precise probability distribution \mathbf{p} and another precise probability \mathbf{q} , or between a precise probability and the nearest or farthest point in some set of imprecise probabilities, in terms of the gain or loss that a decision maker experiences as a result of that difference. For example, \mathbf{q} might be a prior probability distribution over some set of events, which is later updated to a posterior distribution \mathbf{p} based on new information, and the magnitude

of the difference between \mathbf{p} and \mathbf{q} might determine the quantity or value of that information for purposes of signal transmission or decision making. Or, \mathbf{p} might be the precise probability of a decision maker who has the opportunity to bet or trade against an opponent whose beliefs are described by a precise probability \mathbf{q} or by a set of imprecise probabilities \mathcal{Q} , in which case the decision maker can obtain a greater expected payoff or expected utility the farther that \mathbf{p} is from \mathbf{q} or from the nearest point in \mathcal{Q} . Or, the decision maker's probability \mathbf{p} might itself be imprecise, known only to lie within some set \mathcal{P} , and it might be of interest to find the distribution that is nearest to the center of \mathcal{P} in the sense of minimizing the maximum loss that the decision maker could suffer by acting upon the wrong probability.

The considerable literature on this topic includes (at least) three distinct but intertwined strands: scoring rules, entropy, and decision analysis. Scoring rules are reward functions for eliciting and evaluating probability forecasts, and the expected score associated with a forecast can be interpreted as a measure of the value of the forecaster's information. Entropy is a measure of the channel capacity required to communicate a stream of signals generated by a stationary process, and relative entropy measures the reduction in channel capacity that is possible when new information yields an updated signal distribution. Decision analysis provides a general framework for measuring information in terms of gains in expected utility, as well for determining how to optimally use information to choose portfolios of financial assets.

These information-theoretic tools have been used for many decades, but new applications and theoretical developments have emerged during the last few years on several fronts, including experimental economics, robust Bayesian statistics, and financial engineering. The objective of this paper is to add to this recent stream of interdisciplinary literature by broadening the concept of a scoring rule to include a not-

necessarily-uniform baseline distribution and to show that this leads immediately to tight connections with some well-known measures of divergence (relative entropy) as well as with models of utility maximization in markets under uncertainty. In the setting where some probabilities are imprecise, we focus on the problem in which \mathbf{p} is outside the set \mathcal{Q} and the quantity of interest is the divergence between \mathbf{p} and its nearest neighbor in \mathcal{Q} . More details and proofs of the main results are given in Jose et al. (2007)

2 Scoring rules

Scoring rules are reward functions for eliciting and evaluating probabilities, and they have played an important role in the foundations of subjective probability theory (de Finetti 1937 & 1974, Good 1952, Winkler 1967 & 1996, Savage 1971, Lindley 1982) as well as practical applications such as incentive schemes for paying weather forecasters (Brier 1950) and subjects in economic experiments (Selten 1998) and for evaluating the quality of forecasts used in risk analysis (Cooke 1991). Consider an individual (the “forecaster”) who is asked to assess a probability distribution over a set of n mutually exclusive and collectively exhaustive events. Let \mathbf{p} denote the forecaster’s true distribution, let $\mathbf{r} = (r_1, \dots, r_n)$ denote her reported distribution (if different from \mathbf{p}), and let \mathbf{e}_i denote the probability distribution that assigns probability 1 to event i and zero to all other events, i.e., the indicator vector for event i . A scoring rule is conventionally expressed as a function $S(\mathbf{r}, \mathbf{p})$, linear in its second argument, such that the score obtained if event i occurs is $S(\mathbf{r}, \mathbf{e}_i)$, and the forecaster’s *expected* score for reporting \mathbf{r} when her true distribution is \mathbf{p} is $S(\mathbf{r}, \mathbf{p}) = \sum_i p_i S(\mathbf{r}, \mathbf{e}_i)$. It is assumed that the forecaster’s objective is to maximize her expected score, which means that either she is risk neutral and $S(\mathbf{r}, \mathbf{e}_i)$ is measured in units of money or else she is non-risk neutral and $S(\mathbf{r}, \mathbf{e}_i)$ is measured in units of utility.

The scoring rule is defined to be [*strictly*] *proper* if it encourages honest reporting in the sense that $S(\mathbf{p}, \mathbf{p}) \geq S(\mathbf{r}, \mathbf{p})$ for every \mathbf{r} and \mathbf{p} [with equality only when $\mathbf{r} = \mathbf{p}$], so that the forecaster whose true distribution is \mathbf{p} maximizes her expected score by truthfully reporting \mathbf{p} rather than some other distribution. The forecaster’s optimal expected score that is obtained when her distribution is \mathbf{p} will be denoted by merely suppressing the first argument: $S(\mathbf{p}) \equiv S(\mathbf{p}, \mathbf{p})$. A proper scoring rule has a canonical representation in terms of its optimal-expected-score function, as noted by McCarthy (1956) and further elaborated by Hendrickson and Buehler (1971) and Savage (1971). In particular, if $S(\cdot)$ is a differentiable function, then

$S(\cdot, \cdot)$ is uniquely determined by the formula

$$S(\mathbf{r}, \mathbf{p}) = S(\mathbf{r}) + \nabla S(\mathbf{r}) \cdot (\mathbf{p} - \mathbf{r}). \quad (1)$$

where $\nabla S(\mathbf{r})$ denotes the gradient of $S(\cdot)$ evaluated at \mathbf{r} , and conversely every function S that is [strictly] convex and differentiable uniquely defines a [strictly] proper scoring rule. Written in this form, the expected score yielded by a proper scoring rule is seen to be closely related to a particular measure of divergence between probability distributions that is known as a *Bregman divergence* (Brègman 1967), a connection that has been discussed by Grünwald and Dawid (2004), Dawid (2006), and Gneiting and Raftery (2007). Any strictly convex function F defines a Brègman divergence $B_F(\mathbf{p}||\mathbf{r})$ as follows:

$$B_F(\mathbf{p}||\mathbf{r}) = F(\mathbf{p}) - F(\mathbf{r}) - \nabla F(\mathbf{r}) \cdot (\mathbf{p} - \mathbf{r}).$$

Letting $F(\mathbf{p}) = S(\mathbf{p})$, it follows that for any strictly proper scoring rule, the function $S(\mathbf{p}) - S(\mathbf{r}, \mathbf{p})$, which represents the forecaster’s expected *loss* for reporting \mathbf{r} when her true distribution is \mathbf{p} , is a Brègman divergence, and vice versa. Thus, there is a one-to-one correspondence between strictly proper scoring rules and Brègman divergences.

The literature of scoring rules has mainly focused on a few strictly proper rules with particularly convenient parametric forms, axiomatic representations, and/or geometrical interpretations, namely the *quadratic*, *logarithmic*, and *spherical* scoring rules. The quadratic rule (a.k.a. “Brier score”) is $S(\mathbf{p}, \mathbf{e}_i) = -(\|\mathbf{e}_i - \mathbf{p}\|_2)^2$. Thus, under the quadratic rule, the forecast \mathbf{p} is treated as an estimate of the indicator vector of the uncertain event \mathbf{e}_i , and the forecaster is ultimately penalized in proportion to the squared Euclidean distance between \mathbf{p} and the realized value of \mathbf{e}_i , in the tradition of least squares estimation. The logarithmic scoring rule is $S(\mathbf{p}, \mathbf{e}_i) = \ln(p_i)$, whose optimal expected score function is the negative entropy of the forecaster’s true distribution, an issue to which we return below. (Some prescient comments on the potential connection between scoring rules and entropy were made by Good (1971).) The spherical scoring rule is $S(\mathbf{p}, \mathbf{e}_i) = p_i / \|\mathbf{p}\|_2$, and it is obtained by letting the set of feasible score vectors be the simplest strictly convex object in \mathbb{R}^n , namely the unit sphere.

The quadratic and spherical rules can be generalized into parametric families by replacing the 2-norm with the vector β -norm, $\|\mathbf{p}\|_\beta \equiv \left(\sum_{j=1}^n p_j^\beta\right)^{1/\beta}$. The generalized spherical rule is the *pseudospherical scoring rule*, $p_i / (\|\mathbf{p}\|_\beta)^{\beta-1}$, which was first proposed by Good (1971). The generalized quadratic rule is the *power scoring rule*, $\beta p_i^{\beta-1} - (\beta-1) \left(\|\mathbf{p}\|_\beta\right)^\beta$. Written in this

conventional fashion, these families of rules are well-defined and proper only for $\beta > 1$ and the corresponding optimal-expected-score functions that generate them via McCarthy's formula are simply $\left(\|\mathbf{p}\|_\beta\right)^\beta$ and $\|\mathbf{p}\|_\beta$, respectively. The logarithmic scoring rule is the limiting case of both the pseudospherical and power scores as $\beta \rightarrow 1$, but otherwise the two families do not intersect.

3 Weighted score rules and divergence measures

A key property of the aforementioned scoring rules is that they treat events *symmetrically* in the sense that if $p_i = [>] p_j$, then the score in event i is equal to [greater than] the score in event j , regardless of the descriptions of the events, and the forecaster's expected score is smallest when \mathbf{p} is the uniform distribution. Thus, they implicitly reward the forecaster in proportion to some measure of the difference of \mathbf{p} from a uniform distribution. However, in most real (and even hypothetical) applications, the relevant reference point is not a uniform distribution. For example, in weather forecasting the events that are of interest are often known to have widely varying a priori probabilities, and "baseline" values for those probabilities, upon which the forecaster is supposed to improve, are obtainable from historical records (Winkler 1994) or alternative forecasting models. In predicting the outcomes of sporting events or movements of financial markets, there are public betting lines or posted prices for contingent claims that implicitly assign probabilities to events. Therefore, we propose that scoring rules should be generalized so as to reward the forecaster in proportion to some measure of the difference between \mathbf{p} and an appropriate baseline distribution \mathbf{q} . Such a scoring rule will be henceforth referred to as *weighted scoring rule*; it will be expressed as a function of three arguments, $S(\mathbf{r}, \mathbf{p} \parallel \mathbf{q})$, and its associated optimal expected score will be expressed as a function of two arguments, $S(\mathbf{p} \parallel \mathbf{q})$.

There are various functional forms through which the dependence of the score on the baseline distribution could be modeled, and the one we that we find most compelling, for both practical and theoretical reasons, is that for fixed \mathbf{p} and \mathbf{q} the score in state i should depend on the ratio p_i/q_i , so that if $p_i/q_i = [>] p_j/q_j$, then the score in event i should be equal to [greater than] the score in event j . One simple rationale for this desideratum is that when bets may be placed on outcomes of events, *relative* rather than absolute differences in probabilities are what matter, insofar as a \$1 bet on state i has an expected payoff of $\$p_i/q_i$ when the bettor's probability is p_i and the posted odds are

based on q_i . Another rationale can be illustrated by a simple example: suppose that the state space consists of 4 states formed by the Cartesian product of two binary events E and F , and suppose it happens that the forecaster and client both agree on the probability of F and they also agree that E and F are probabilistically independent. Then it seems reasonable that the forecaster's payment should depend only on the outcome of E , not F , and this requires the payoff in each of the four states to depend only on the ratio of p to q , which is the relative change in the evaluation of the probability of E .

The measurement of differences between probabilities in terms of ratios has a long history in statistics and information theory. It was noted above that under a strictly proper scoring rule, the forecaster's expected *loss* for reporting a distribution \mathbf{r} that is other than her true distribution \mathbf{p} is a particular kind of divergence between \mathbf{r} and \mathbf{p} , namely a Brègman divergence. Under a weighted strictly proper scoring rule that bases the score on the ratio p_i/q_i the forecaster's expected *gain* for possessing a distribution \mathbf{p} that differs from \mathbf{q} is a second kind of divergence, which is not a Brègman divergence. Rather, it turns out to be a special case (or a simple transformation) of another kind of generalized divergence known as an f -divergence (Csiszár 1967). If f is a strictly convex function, the corresponding f -divergence is defined as

$$D_f(\mathbf{p} \parallel \mathbf{q}) = E_{\mathbf{p}}[f(\mathbf{p}/\mathbf{q})]. \quad (2)$$

Divergences of this general form have been widely used in statistics for many years as "utility-free" measures of the value of the information - e.g., Goel (1983) uses f -divergence to define a "conditional amount of sample information" for measuring prior-to-posterior information gains in Bayesian hierarchical models. More recently it has been recognized that f -divergences are interpretable as measures of expected utility gains that are available to decision makers who have opportunities to bet against less-well-informed opponents or to invest in financial markets, as will be more fully discussed in later sections of this paper.

When the ratio p_i/q_i is substituted for p_i in the pseudospherical and power scoring rules, and they are affinely transformed so as to yield scores of zero when $\mathbf{p} = \mathbf{q}$, we obtain the *weighted pseudospherical score*, denoted S_β^S , and the *weighted power score*, denoted

by $S_\beta^{\mathbf{P}}$, with the following parametric forms:

$$S_\beta^{\mathbf{S}}(\mathbf{p}, \mathbf{e}_i \| \mathbf{q}) \equiv \frac{1}{\beta - 1} \left(\left(\frac{p_i/q_i}{(E_{\mathbf{p}}[(\mathbf{p}/\mathbf{q})^{\beta-1}]^{1/\beta})} \right)^{\beta-1} - 1 \right), \quad (3)$$

$$S_\beta^{\mathbf{P}}(\mathbf{p}, \mathbf{e}_i \| \mathbf{q}) \equiv \frac{(p_i/q_i)^{\beta-1} - 1}{\beta - 1} - \frac{E_{\mathbf{p}}[(\mathbf{p}/\mathbf{q})^{\beta-1}] - 1}{\beta}. \quad (4)$$

Note that for any fixed values of \mathbf{p} , \mathbf{q} , and β , the pseudospherical score vector $(S_\beta^{\mathbf{S}}(\mathbf{p}, \mathbf{e}_1 \| \mathbf{q}), \dots, S_\beta^{\mathbf{S}}(\mathbf{p}, \mathbf{e}_n \| \mathbf{q}))$ is a positive affine transformation of the power score vector $(S_\beta^{\mathbf{P}}(\mathbf{p}, \mathbf{e}_1 \| \mathbf{q}), \dots, S_\beta^{\mathbf{P}}(\mathbf{p}, \mathbf{e}_n \| \mathbf{q}))$, since both vectors are affine transformations of $(\mathbf{p}/\mathbf{q})^{\beta-1}$, although the origins and scale factors of the transformations vary with \mathbf{p} , \mathbf{q} , and β . Thus, although the two rules yield different expected payoffs as a function of \mathbf{p} (for the same \mathbf{q} and β), and they create different incentives for information-gathering and different penalties for dishonest reporting, they nevertheless present the same relative risk profile to a truthful forecaster whose \mathbf{p} is already fixed. At $\beta = 1$ both rules converge to the weighted logarithmic score $\ln(p_i/q_i)$. At $\beta = 2$, weighted forms of the quadratic and spherical scoring rules are obtained. The cases $\beta = 0$ and $\beta = \frac{1}{2}$ have not received much (if any) attention in the antecedent literature, but it will be shown later that $\beta = 0$ corresponds to a decision model involving exponential utility, which is the utility function most commonly used in applied decision analysis, while $\beta = \frac{1}{2}$ arises from a decision model involving reciprocal utility, which has some appealing symmetry properties and is closely related to the Hellinger distance between \mathbf{p} and \mathbf{q} . These special cases will be further explored in the next two sections.

The corresponding optimal-expected-score functions for the two families of weighted scoring rules are:

$$S_\beta^{\mathbf{S}}(\mathbf{p} \| \mathbf{q}) = \frac{(E_{\mathbf{p}}[(\mathbf{p}/\mathbf{q})^{\beta-1}]^{1/\beta} - 1}{\beta - 1}, \quad (5)$$

$$S_\beta^{\mathbf{P}}(\mathbf{p} \| \mathbf{q}) = \frac{E_{\mathbf{p}}[(\mathbf{p}/\mathbf{q})^{\beta-1}] - 1}{\beta(\beta - 1)}, \quad (6)$$

and one is a monotonically increasing function of the other for any fixed β . Our first result is to point out that these expected score functions correspond exactly to two parametric families of generalized *divergence* (cross-entropy) between probability distributions. In particular the weighted power expected score $S_\beta^{\mathbf{P}}(\mathbf{p} \| \mathbf{q})$ is precisely the *directed divergence of order β between \mathbf{p} and \mathbf{q}* proposed by Havrda and

Chavráč (1967), variants of which have been discussed by Rathie and Kannappan (1972), Cressie and Read (1984), and Haussler and Opper (1997). Cressie and Read refer to this quantity as the *power divergence*, which we shall also do here.

The weighted pseudospherical score $S_\beta^{\mathbf{S}}(\mathbf{p} \| \mathbf{q})$ is the cross-entropy measure that arises from a generalized entropy introduced by Arimoto (1971) and further elaborated by Sharma and Mittal (1975), Boekee and Van der Lubbe (1980) and Lavenda and Dunning-Davies (2003). Arimoto's generalized entropy of order β is defined for $\beta > 0$ by $\beta/(\beta - 1) (E_{\mathbf{p}}[\mathbf{p}^{\beta-1}]^{1/\beta} - 1)$. The factor of β in the numerator plays no essential role when β is restricted to be positive, and without it the measure is actually valid for all real β , and when $\mathbf{p}^{\beta-1}$ is replaced by $(\mathbf{p}/\mathbf{q})^{\beta-1}$ so as to define a cross-entropy, the weighted pseudospherical expected score is obtained. It is therefore appropriate to refer to the latter quantity as the *pseudospherical divergence of order β between \mathbf{p} and \mathbf{q}* . Both of these generalized divergences reduce to the Kullback-Leibler divergence $E_{\mathbf{p}}[\ln(\mathbf{p}/\mathbf{q})]$ at $\beta = 1$, and for other special cases of β they are related to two other well known divergences, namely the Chi-square divergence $\chi^2(\mathbf{q} \| \mathbf{p}) = E_{\mathbf{p}}[\mathbf{p}/\mathbf{q}] - 1$ and the Hellinger distance $D_H(\mathbf{p} \| \mathbf{q}) \equiv \left(\sum_{j=1}^n (\sqrt{p_j} - \sqrt{q_j})^2 \right)^{1/2}$ as shown in the following table:

Table 1. Power & pseudospherical divergences

β	$S_\beta^{\mathbf{P}}(\mathbf{p} \ \mathbf{q})$	$S_\beta^{\mathbf{S}}(\mathbf{p} \ \mathbf{q})$
-1	$\frac{1}{2} \chi^2(\mathbf{q} \ \mathbf{p})$	$\frac{1}{2} (1 - (\chi^2(\mathbf{q} \ \mathbf{p}) + 1)^{-1})$
0	$D_{KL}(\mathbf{q} \ \mathbf{p})$	$1 - \exp(-D_{KL}(\mathbf{q} \ \mathbf{p}))$
$\frac{1}{2}$	$2D_H(\mathbf{p} \ \mathbf{q})^2$	$2 \left(1 - \left(1 - \frac{1}{2} D_H(\mathbf{p} \ \mathbf{q})^2 \right)^2 \right)$
1	$D_{KL}(\mathbf{p} \ \mathbf{q})$	$D_{KL}(\mathbf{p} \ \mathbf{q})$
2	$\frac{1}{2} \chi^2(\mathbf{p} \ \mathbf{q})$	$\sqrt{\chi^2(\mathbf{p} \ \mathbf{q}) + 1} - 1$

Note that the power divergence is symmetric around $\beta = \frac{1}{2}$ in the sense that $S_\beta^{\mathbf{P}}(\mathbf{p} \| \mathbf{q}) = S_{1-\beta}^{\mathbf{P}}(\mathbf{q} \| \mathbf{p})$, i.e., the roles of \mathbf{p} and \mathbf{q} are merely reversed when β is replaced by $1 - \beta$.

4 Decision models and information measures

Our second result is to show that the same two families of generalized divergence arise naturally as the solutions of two canonical expected-utility maximization problems, involving the most widely-used parametric family of utility functions, in which a risk averse decision maker with subjective probability dis-

tribution \mathbf{p} bets against a non-strategic risk-neutral opponent with distribution \mathbf{q} , or equivalently, invests in a contingent claims market where prices are determined by taking expectations with respect to \mathbf{q} . A contingent claim is a claim to monetary payments that are contingent on states of the world, and it can be represented as an n -vector of payoffs \mathbf{y} that has some market price $p(\mathbf{y})$ at which it can be purchased in arbitrary positive multiples. (In a financial market, the relevant states of the world might be possible values of a stock price or stock index on a particular future date, and a contingent claim might be a share of stock or an option to buy a share of that stock at a pre-specified strike price.) A decision maker who buys α units of \mathbf{y} at its market price receives a net payoff of $\alpha(y_i - p(\mathbf{y}))$ in state i , hence the vector $\alpha(\mathbf{y} - p(\mathbf{y})\mathbf{1})$ is a feasible net payoff vector for the decision maker for all positive α . The market is *complete* if every contingent claim has a unique price at which it can be both bought *and* sold, in which case $\alpha(\mathbf{y} - p(\mathbf{y})\mathbf{1})$ is a feasible payoff vector for all real α , positive or negative. If the market prices are also *arbitrage-free* (“coherent”), then there exists a unique probability distribution \mathbf{q} that prices all contingent claims according to their expected payoffs, so that $p(\mathbf{y}) = E_{\mathbf{q}}[\mathbf{y}]$ for all $\mathbf{y} \in \mathbb{R}^n$, and any $\mathbf{x} \in \mathbb{R}^n$ that satisfies $E_{\mathbf{q}}[\mathbf{x}] = 0$ is a feasible net payoff vector. In Bayesian theory this existence result is known as de Finetti’s “fundamental theorem of probability,” with $p(\mathbf{y})$ referred to as the “prevision” of \mathbf{y} , and in finance theory it is known as the “fundamental theorem of asset pricing,” with \mathbf{q} referred to as a “risk neutral distribution” because assets are priced “as if” by a risk neutral opponent whose probability distribution is \mathbf{q} .

In the first canonical problem (“S”), there is a single time period in which consumption occurs and the decision maker has a single-attribute vNM utility function $u(x)$. The decision maker’s optimal expected utility, denoted $U^{\mathbf{S}}(\mathbf{p}||\mathbf{q})$, is determined by:

$$\begin{aligned} \textbf{Problem S} & : & (7) \\ U^{\mathbf{S}}(\mathbf{p}||\mathbf{q}) & \equiv \max_{\mathbf{x} \in \mathbb{R}^n} E_{\mathbf{p}}[u(\mathbf{x})] \quad \text{s.t. } E_{\mathbf{q}}[\mathbf{x}] = 0, \end{aligned}$$

where $u(\mathbf{x}) \equiv (u(x_1), \dots, u(x_n))$ denotes the vector of utilities that u yields when applied to \mathbf{x} . In the second problem (“P”), there are two periods in which consumption occurs and the decision maker with probability distribution \mathbf{p} has a quasilinear vNM utility function $u(a, b) = a + u(b)$ where a is money consumed at time 0 and b is money consumed at time 1. The decision maker’s objective is to choose a vector \mathbf{x} of time-1 payoffs to be purchased from time-0 funds at market prices so as to maximize the expected utility of consumption in both periods. The time-0

cost of purchasing \mathbf{x} is $E_{\mathbf{q}}[\mathbf{x}]$, so the optimal expected utility, denoted $U^{\mathbf{P}}(\mathbf{p}||\mathbf{q})$, is the solution of:

$$\begin{aligned} \textbf{Problem P} & : & (8) \\ U^{\mathbf{P}}(\mathbf{p}||\mathbf{q}) & \equiv \max_{\mathbf{x} \in \mathbb{R}^n} E_{\mathbf{p}}[u(\mathbf{x})] - E_{\mathbf{q}}[\mathbf{x}]. \end{aligned}$$

Next, let u be a utility function from the general exponential/logarithmic/power family, which will be parameterized here as:

$$u_{\beta}(x) \equiv \frac{1}{\beta - 1} ((1 + \beta x)^{(\beta - 1)/\beta} - 1) \quad (9)$$

for $\beta x > -1$. This parameterization has two key properties. First, $u_{\beta}(0) = 0$ and $u'_{\beta}(0) = 1$, so that for any β the marginal rate of substitution between time-0 consumption and time-1 consumption is unity at $x = 0$ for the decision maker in Problem P. Second, the corresponding *risk tolerance function* $\tau_{\beta}(x)$, which is the reciprocal of the Pratt-Arrow risk aversion measure, is the following linear function of wealth: $\tau_{\beta}(x) \equiv -u'_{\beta}(x)/u''_{\beta}(x) = 1 + \beta x$. Thus, the risk tolerance as well as the marginal utility is normalized to a value of 1 at $x = 0$, and β is the coefficient of risk tolerance, i.e., the increase in risk tolerance per unit of increase in wealth. The linear-risk-tolerance utility functions are also known as hyperbolic-absolute-risk-aversion (HARA) utility functions in the literature of financial economics, although parameterizing them in terms of their risk tolerance coefficients rather than their risk aversion coefficients is more useful for our purposes. Some important special cases of u_{β} are given in Table 2:

Table 2. Linear-risk-tolerance utility functions

β	$u_{\beta}(x)$	Functional form
-1	$u_{-1}(x) = -\frac{1}{2}((1-x)^2 - 1)$	Quadratic
0	$u_0(x) = 1 - \exp(-x)$	Exponential
$\frac{1}{2}$	$u_{1/2}(x) = 2 \left(1 - \frac{1}{1+x/2}\right)$	Reciprocal
1	$u_1(x) = \ln(1+x)$	Logarithmic
2	$u_2(x) = \sqrt{1+2x} - 1$	Square-root

The utility functions $\{u_{\beta}\}$ also exhibit a symmetry around $\beta = \frac{1}{2}$, namely that $u_{1-\beta}(x) = -u_{\beta}(-x)$, or equivalently $u_{\beta}(-u_{1-\beta}(-x)) = x$. In other words, the graph of $u_{1-\beta}$ is obtained from the graph of u_{β} by merely reflecting it around the line $y = -x$. Note that the power (exponent) in u_{β} is the term $(\beta - 1)/\beta$, which has the property that $((\beta - 1)/\beta)^{-1} = ((1 - \beta) - 1)/(1 - \beta)$, so that swapping β for $1 - \beta$ results in another power utility function whose power is the reciprocal of the original. Thus, up to affine scaling, the reciprocal utility function ($\beta = \frac{1}{2}$) is self-symmetric, the exponential and logarithmic utility

functions ($\beta = 0$ and $\beta = 1$) are symmetric to each other, and the power utility function with exponent δ is symmetric to the power utility function with exponent $1/\delta$ for any positive or negative δ other than 0 or 1.

Let $\mathbf{x}_\beta^{\mathbf{S}}(\mathbf{p}|\mathbf{q})$ and $\mathbf{x}_\beta^{\mathbf{P}}(\mathbf{p}|\mathbf{q})$ denote the solutions of Problems **S** and **P** when $u = u_\beta$, with i^{th} elements $x_{\beta,i}^{\mathbf{S}}(\mathbf{p}|\mathbf{q})$ and $x_{\beta,i}^{\mathbf{P}}(\mathbf{p}|\mathbf{q})$, respectively, and let $U_\beta^{\mathbf{S}}(\mathbf{p}|\mathbf{q})$ and $U_\beta^{\mathbf{P}}(\mathbf{p}|\mathbf{q})$ denote their corresponding expected utilities. In these terms, we have:

THEOREM 1:

- (a) $S_\beta^{\mathbf{S}}(\mathbf{p}, \mathbf{e}_i|\mathbf{q}) = u_\beta(x_{\beta,i}^{\mathbf{S}}(\mathbf{p}|\mathbf{q}))$,
and $S_\beta^{\mathbf{S}}(\mathbf{p}|\mathbf{q}) = U_\beta^{\mathbf{S}}(\mathbf{p}|\mathbf{q})$
- (b) $S_\beta^{\mathbf{P}}(\mathbf{p}, \mathbf{e}_i|\mathbf{q}) = u_\beta(x_{\beta,i}^{\mathbf{P}}(\mathbf{p}|\mathbf{q})) - E_{\mathbf{q}}[\mathbf{x}_\beta^{\mathbf{P}}(\mathbf{p}|\mathbf{q})]$,
and $S_\beta^{\mathbf{P}}(\mathbf{p}|\mathbf{q}) = U_\beta^{\mathbf{P}}(\mathbf{p}|\mathbf{q})$
- (c) $S_\beta^{\mathbf{P}}(\mathbf{p}|\mathbf{q}) \geq S_\beta^{\mathbf{S}}(\mathbf{p}|\mathbf{q})$ for all \mathbf{p} , \mathbf{q} , and β .

Thus, the statewise utility gains to the decision maker under problems **S** and **P** are precisely the pseudospherical and power scores for the same \mathbf{p} , \mathbf{q} , and β , and the expected utilities are the corresponding divergences.

5 Utility/entropy duality in incomplete markets

We now extend the preceding results to a setting in which the decision maker's risk neutral betting opponent has imprecise probabilities, which is equivalent to an incomplete market where a contingent claim may have a "bid-ask spread" rather than a single price at which it can be both bought and sold. The bid-ask spreads generally do not suffice to determine a unique risk neutral distribution; rather, they only determine a convex set \mathcal{Q} of risk-neutral distributions such that \mathbf{x} is a feasible net payoff vector for the decision maker if and only if $E_{\mathbf{q}}[\mathbf{x}] \leq 0$ for all $\mathbf{q} \in \mathcal{Q}$. (The payoff to the opponent is $-\mathbf{x}$, hence the constraint $E_{\mathbf{q}}[\mathbf{x}] \leq 0$ for all $\mathbf{q} \in \mathcal{Q}$ means that the opponent with imprecise probabilities \mathcal{Q} will accept only those bets yielding non-negative expected payoffs for all $\mathbf{q} \in \mathcal{Q}$.) The problem of expected-utility maximization in incomplete markets has been widely studied in the mathematical finance literature in recent years, and it has been shown that there is a duality relationship between maximization of expected utility and minimization of an appropriate measure of relative entropy or divergence (e.g., Frittelli 2000, Rouge and El Karoui 2000, Goll and Rüschenendorf 2001, Delbaen et al. 2002, Slomczyński and Zastawniak 2004, Ilhan et al. 2004, Samperi 2005). Most of this lit-

erature has focused on the case of exponential utility, for which the dual problem is the minimization of the reverse KL divergence $D_{KL}(\mathbf{q}, \mathbf{p})$, as well as on issues that arise in multi-period or continuous-time markets. In this section we will show that in a single-period or two-period market, the duality relationship applies to the entire spectrum of linear-risk-tolerance utility and pseudospherical divergence or power divergence.

An incomplete, single-period market can be parameterized in either of two ways. One is in terms of an $m \times n$ matrix \mathbf{A} whose rows are feasible net payoff vectors, i.e., $\mathbf{A} = \{a_{ij}\}$ where a_{ij} is the net payoff that the decision maker receives in the j^{th} state of the world for purchasing one unit of the i^{th} contingent claim at its asking price. (It suffices to consider only purchases at asking prices, rather than sales at bid prices, since a bid price of p for a contingent claim \mathbf{y} is equivalent to an asking price of $-p$ for $-\mathbf{y}$.) Alternatively, the market can be parameterized in terms of a $k \times n$ matrix \mathbf{Q} whose rows are risk neutral probability distributions that support the contingent claim prices, i.e., $\mathbf{Q} = \{q_{ij}\}$ where q_{ij} is the probability of the j^{th} state of the world under the i^{th} risk neutral distribution. The rows of \mathbf{Q} are the extremal risk-neutral probability distributions assigning non-positive expectation to all the rows of \mathbf{A} , i.e., the rows of $-\mathbf{Q}$ are the dual cone of the rows of \mathbf{A} . The second parameterization will be adopted here, in terms of which \mathcal{Q} is the convex hull of the rows of \mathbf{Q} , so that a generic element of \mathcal{Q} can be expressed as $\mathbf{q} = \mathbf{z}^T \mathbf{Q}$ where \mathbf{z} is an element of Δ^k , the unit simplex in \mathbb{R}^k , and the feasibility requirement that $E_{\mathbf{q}}[\mathbf{x}] \leq 0$ for all $\mathbf{q} \in \mathcal{Q}$ can be expressed as $\mathbf{Q}\mathbf{x} \leq \mathbf{0}$.

In the incomplete-market generalization of Problem **S**, the problem of finding the maximum expected utility, which will be denoted as $U_\beta^{\mathbf{S}}(\mathbf{p}|\mathbf{Q})$, is dual to the problem of finding the minimum pseudospherical divergence of order β between \mathbf{p} and all \mathbf{q} in the convex hull of the rows of \mathbf{Q} , which will be denoted as $S_\beta^{\mathbf{S}}(\mathbf{p}|\mathbf{Q})$:

Primal Problem S :

$$U_\beta^{\mathbf{S}}(\mathbf{p}|\mathbf{Q}) \equiv \max_{\mathbf{x} \in \mathbb{R}^n} E_{\mathbf{p}}[u_\beta(\mathbf{x})], \mathbf{Q}\mathbf{x} \leq \mathbf{0}$$

Dual Problem S :

$$S_\beta^{\mathbf{S}}(\mathbf{p}|\mathbf{Q}) \equiv \min_{\mathbf{z} \in \Delta^k} S_\beta^{\mathbf{S}}(\mathbf{p}|\mathbf{z}^T \mathbf{Q}).$$

In the incomplete-market generalization of Problem **P**, the decision maker's objective is to determine an amount w to be spent at time 0 to finance consumption in period 1. For the period-1 payoff vector \mathbf{x} that the decision maker wishes to purchase, the risk-neutral expected value of \mathbf{x} needs to be less than or equal to w for all the extremal risk neutral distributions. The corresponding primal and dual problems

are:

$$\text{Primal Problem } \mathbf{P} : U_{\beta}^{\mathbf{P}}(\mathbf{p} \parallel \mathbf{Q}) \equiv \max_{\mathbf{x} \in \mathbb{R}^n} E_{\mathbf{p}}[u_{\beta}(\mathbf{x})] - w, \mathbf{Q}\mathbf{x} \leq w\mathbf{1}$$

$$\text{Dual Problem } \mathbf{P} : S_{\beta}^{\mathbf{P}}(\mathbf{p} \parallel \mathbf{Q}) \equiv \min_{\mathbf{z} \in \Delta^k} S_{\beta}^{\mathbf{P}}(\mathbf{p} \parallel \mathbf{z}^T \mathbf{Q}).$$

The special case $\beta = 1$ corresponds to logarithmic utility in the primal problem and KL divergence in the dual problem, while $\beta = 0$ corresponds to exponential utility in the primal problem and reverse KL divergence in the dual problem, and the cases $\beta = 1/2$ and $\beta = \pm 2$ are related to the squared Hellinger distance and the Chi-square divergence as shown in Table 1. These duality relationships are formalized in:

THEOREM 2:

(a) In an incomplete, single-period market, maximization of expected linear-risk-tolerance utility with risk tolerance coefficient β (Primal Problem **S**) is equivalent to minimization of the pseudospherical divergence of order β between the decision maker’s subjective distribution \mathbf{p} and a risk neutral distribution \mathbf{q} consistent with contingent claim prices (Dual Problem **S**). Their optimal objective values are the same and the optimal values of the decision variables in one problem are equal to the normalized optimal values of the Lagrange multipliers in the other.

(b) In an incomplete, two-period market, maximization of quasi-linear expected linear-risk-tolerance utility with second-period risk tolerance coefficient β (Primal Problem **P**) is equivalent to minimization of the power divergence of order β between the decision maker’s subjective distribution \mathbf{p} and a risk neutral distribution \mathbf{q} consistent with contingent claim prices (Dual Problem **P**). Their optimal objective values are the same and the optimal values of the decision variables in one problem are equal to the normalized optimal values of the Lagrange multipliers in the other.

Note that because the pseudospherical divergence is a monotonic transformation of the power divergence, the distribution \mathbf{q} ($= \mathbf{z}^T \mathbf{Q}$) that solves Dual Problem **S** is the same one that solves Dual Problem **P**, although the objective values and the primal payoff vectors are generally different. The geometry of the dual solutions is illustrated in Figure 1.

Grünwald and Dawid (2004) have explored duality relationships among strictly proper scoring rules, generalized entropies and divergences, and expected-utility-maximization (or in their terms, expected-loss-minimization) in the context of robust Bayesian inference, where the decision maker does not know the true probability distribution and her opponent is “Nature” who chooses the true distribution \mathbf{p} from some

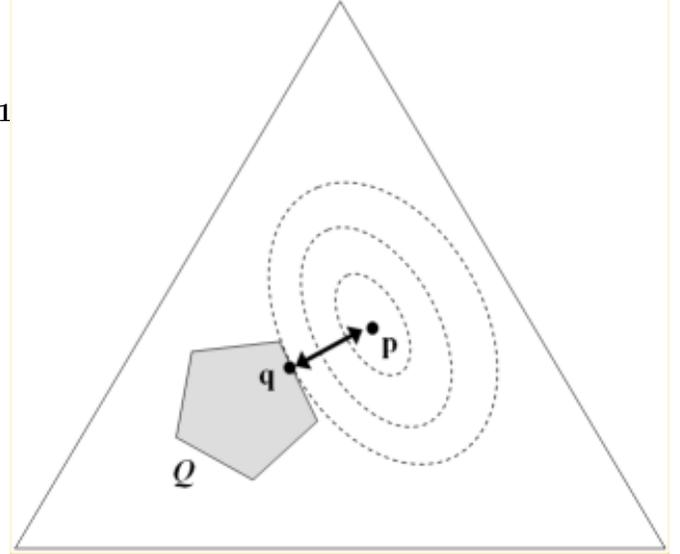


Figure 1: Geometry of minimizing the divergence between \mathbf{p} and the nearest element of \mathcal{Q} ($n = 3$)

convex set \mathcal{P} , such as the set of distributions satisfying a mean-value constraint. The robust Bayes problem for the decision maker is to determine the distribution \mathbf{r} that minimizes her maximum expected loss over all $\mathbf{p} \in \mathcal{P}$, where the expected loss (in our terms) is the negative expected score $-S(\mathbf{r}, \mathbf{p})$. Grünwald and Dawid show that the optimal-expected-loss function, $-S(\mathbf{p})$, is interpretable as a generalized entropy, and minimizing the maximum expected loss is equivalent to maximizing this entropy on the set \mathcal{P} . This scoring-rule entropy uniquely determines a corresponding Brègman divergence $B_S(\mathbf{p} \parallel \mathbf{r}) \equiv S(\mathbf{p}) - S(\mathbf{r}, \mathbf{p})$, as noted earlier, and Grünwald and Dawid go on to show that the distribution \mathbf{r} that minimizes the maximum expected loss on \mathcal{P} is also the distribution that minimizes this divergence with respect to an uninformative “reference” distribution \mathbf{p}_0 at which the entropy $-S(\mathbf{p})$ is maximized. For typical symmetric scoring rules, the reference distribution is the uniform distribution, but any scoring rule entropy can be transformed so as to shift the reference point to any other distribution \mathbf{p}_0^* by the addition of a linear function of \mathbf{p} , namely $S(\mathbf{p}_0^*, \mathbf{p})$. The reference distribution \mathbf{p}_0 in their model therefore plays an analogous role to the baseline distribution \mathbf{q} in our model, insofar as $-S(\mathbf{p} \parallel \mathbf{r})$ is maximized in the uninformative case where $\mathbf{p} = \mathbf{q}$. Grünwald and Dawid also discuss scoring rules for continuous probability distributions drawn from the generalized exponential family, focusing in particular on the logarithmic and quadratic scoring rules.

6 Summary and Conclusions

We have shown that when a risk averse decision maker with a precise probability distribution \mathbf{p} bets against a risk neutral opponent with a convex set \mathcal{Q} of imprecise probabilities, or equivalently invests in an incomplete market for contingent claims where \mathcal{Q} is the set of risk neutral distributions determined by market prices, there is a natural duality between maximizing LRT utility and minimizing pseudospherical or power divergence with the same value of β . In particular, maximization of logarithmic utility ($\beta = 1$) corresponds to finding the distribution \mathbf{q} in \mathcal{Q} that minimizes the KL divergence $D_{KL}(\mathbf{p}||\mathbf{q})$, maximization of exponential utility ($\beta = 0$) corresponds to minimizing the reverse KL divergence $D_{KL}(\mathbf{q}||\mathbf{p})$, and maximization of reciprocal utility ($\beta = \frac{1}{2}$) or square-root utility ($\beta = 2$) correspond to minimization of the Hellinger distance $D_H(\mathbf{p}||\mathbf{q})$ or the Chi-square divergence $\chi^2(\mathbf{p}||\mathbf{q})$, respectively.

References

- [1] ARIMOTO, S. 1971. Information-theoretical considerations on estimation problems. *Infor. Contr.* **19**:181-194.
- [2] BOEKKEE, D. E., J. C. A. VAN DER LUBBE. 1980. The R-norm information measure. *Infor. Contr.* **45**:136-155.
- [3] BRĚGMAN, L. 1967. The relaxation method of finding the common points of convex sets and its application to the solution of problems in convex programming. *USSR Comput. Math. and Math. Phys.*, **7**:200-217.
- [4] BRIER, G. W. 1950. Verification of forecasts expressed in terms of probability. *Mon. Weather Rev.* **78**(1):1-3.
- [5] COOKE, R. 1991. Experts in uncertainty: Opinion and subjective probability in science. Oxford University Press, Oxford.
- [6] CRESSIE, N., T. R. C. READ. 1984. Multinomial goodness of fit. *J. Royal Stat. Soc. B* **46**(3):440-464.
- [7] I. CSISZÁR. 1967 Information type measures of differences of probability distribution and indirect observations. *Studia Math. Hungarica* **2**:299-318
- [8] DAWID, A.P. 2006. The geometry of proper scoring rules. Research Report No. 268, Department of Statistical Science, University College, London. (April 2006)
- [9] DE FINETTI, B. 1937. La Prévision: Ses Lois Logiques, Ses Sources Subjectives. *Ann. Inst. Henri Poincaré* **7**, 1—68. Translation reprinted in H.E. Kyburg and H.E. Smokler, Eds. (1980), *Studies in Subjective Probability*, 2nd ed., Robert Krieger, New York, 53-118.
- [10] DE FINETTI, B. 1974. *Theory of probability, Vol. 1*. Wiley, New York.
- [11] DELBAEN, F., P. GRANDITS, T. RHEINLÄNDER, D. SAMPERI, M. SCHWEIZER, C. STRICKER. 2002. Exponential hedging and entropy penalties. *Math. Finance.* **12**:99-123.
- [12] FRITTELLI, M. 2000. The minimal entropy martingale measure and the valuation problem in incomplete markets. *Math. Finance.* **10**:39-52.
- [13] GNEITING, T., RAFTERY, A. 2007. Strictly proper scoring rules, prediction, and estimation. *J. Am. Stat. Assoc.* **102**:359-378
- [14] GOEL, P. 1983. Information measures and Bayesian hierarchical models. *J. Am. Stat. Assoc.* **78**:408-410.
- [15] GOLL, T., RÜSCHENDORF, L. 2001. Minimax and minimal distance martingale measures and their relationship to portfolio optimization. *Finance Stoch.* **5**:557-581
- [16] GOOD, I.J. 1952 Rational decisions. *J. Royal Statist. Soc. B.* **14**: 107-114.
- [17] GOOD, I.J. 1971. Comment on paper by Buehler. In *Foundations of Statistical Inference* (Godambe and Sprott, eds., Holt, Reinhart, & Winston, Toronto) 337-339.
- [18] GRÜNWARD, P.D., A.P. DAWID. 2004. Game theory, maximum entropy, minimum discrepancy, and robust Bayesian decision theory. *Ann. Statist.* **32**: 1367-1433.
- [19] HAVRDA, J., F. CHAVRÁT. 1967. Quantification method of classification processes: the concept of structural α -entropy. *Kybernetika* **3**:30-35.
- [20] HAUSSLER, D., M. OPPER. 1997. Mutual information, metric entropy, and cumulative relative entropy risk. *Ann. Statist.* **25**: 2451-2492.
- [21] HENDRICKSON, A. D., R. J. BUEHLER 1971. Proper scores for probability forecasters. *Ann. Math. Stat.* **42**:1916-1921.
- [22] ILHAN, A., M. JONSSON, R. SIRCAR. 2004. Portfolio optimization with derivatives and indifference pricing. Working Paper, Princeton University.

- [23] JOSE, V., R. NAU, R. WINKLER. 2007. Scoring Rules, Entropy, and Utility Maximization. Working Paper, Duke University
- [24] LAVENDA, B. H., J. DUNNING-DAVIES. 2003. Qualms concerning Tsallis's condition of pseudo-additivity as a definition of non-extensivity. <http://arxiv.org/abs/cond-mat/0312132>.
- [25] LINDLEY, D. 1982. Scoring rules and the inevitability of probability. *Int. Stat. Rev.* **50**(1):1-26.
- [26] MCCARTHY, J. 1956. Measures of the value of information. *Proc. Nat. Acad. Sci. USA* **42**:654-655.
- [27] RATHIE, P. N., P. KANNAPPAN. 1972. A directed-divergence function of type β . *Infor. Contr.* **20**:38-45.
- [28] ROUGE, R., N. EL KAROUI. 2000. Pricing via utility maximization and entropy. *Math. Finance.* **10**:259-276.
- [29] SAMPERI, D. 2005. Model selection using entropy and geometry: complements to the six-author paper. Working paper, Decision Synergy.
- [30] SAVAGE, L. J. 1971. Elicitation of personal probabilities and expectations. *J. Am. Stat. Assoc.* **66**:783-801.
- [31] SELTEN, R. 1998. Axiomatic characterization of the quadratic scoring rule. *Exper. Econ.* **1**:43-62.
- [32] SHANNON, C. E. 1948. A mathematical theory of communication. *Bell Sys. Tech. J.* **27**:379-423.
- [33] SHARMA, B. D., D. P. MITTAL. 1975. New non-additive measures of entropy for discrete probability distributions *J. Math. Sci.* **10**:28-40.
- [34] SŁOMCZYŃSKI, W., T. ZASTAWIAK. 2004. Utility maximizing entropy and the second law of thermodynamics. *Ann. Prob.* **32**:2261-2285.
- [35] WINKLER, R. L. 1967. The quantification of judgment: some methodological suggestions. *J. Am. Stat. Assoc.* **62**:1105-1120.
- [36] WINKLER, R. L. 1994. Evaluating probabilities: asymmetric scoring rules. *Manage. Sci.* **40**(11):1395-1405.
- [37] WINKLER, R. L. 1996. Scoring rules and the evaluation of probabilities. *Test* **5**(1):1-60.